Adapting to Conflict: Equilibrium Structure and Adaptive Learning in Harmonic Games

Davide Legacci[*] **Panayotis Mertikopoulos**
Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LIG 38000 Grenoble, France
{davide.legacci,panayotis.mertikopoulos}@univ-grenoble-alpes.fr

Bary Pradelski
CNRS, Maison Française d'Oxford
2–10 Norham Road, Oxford, OX2 6SE, United Kingdom
bary.pradelski@cnrs.fr

Abstract

We propose an adaptive, "parameter-agnostic" algorithm for learning in *harmonic games*, a class of finite games characterized by strictly opposed, conflicting interests. In contrast to zero-sum games—which do not always capture conflict, even between two players—harmonic games may exhibit highly non-convex equilibrium configurations. In particular, if the game's ambient dimension is odd, we show that the set of totally mixed Nash equilibria is a (typically non-convex) real algebraic variety of dimension *at least* one, which extends all the way to the boundary of the game's strategy space. Despite the difficulties incurred by this complicated equilibrium structure, we provide an extrapolation-based variant of Follow-the-Regularized-Leader (FTRL) which converges to Nash equilibrium with order-optimal regret guarantees: $\mathcal{O}(1)$ when employed by all players, and $\mathcal{O}(\sqrt{T})$ against arbitrary play (or Byzantine deviations by some of the players).

1 Introduction

Background. Two-player zero-sum games have long served as the canonical model of strategic conflict in non-cooperative game theory, and for good reason: in such games, one player's gain is exactly the other's loss. However—and somewhat counter-intuitively—the zero-sum property does not faithfully capture situations of "strictly opposed interests": there exist simple 2×2 zero-sum games that are strategically equivalent to fully cooperative, *common interest* ones. This complicates the interpretation of min-max games as canonical models of conflict: either one must abandon them as the "worst case" of non-cooperation, or one must concede a fundamental tension between the zero-sum property and the axiomatization of utility theory [45].

A recent twist to this matter was proposed by Candogan et al. [11], who introduced the class of *harmonic games*. Viewed abstractly, this class contains *all* two-player zero-sum games with a totally mixed equilibrium (Matching Pennies, Rock-Paper-Scissors, etc.), while excluding any (non-trivial) zero-sum game which is strategically equivalent to a common interest one. In particular, one of the key results of Candogan et al. [11]—subsequently extended by Abdou et al. [1]—shows that the space of harmonic games is, in a precise sense, *orthogonal* to the space of potential games, and that every game admits a unique decomposition into a potential plus a harmonic component. Consequently, harmonic games have recently come to be viewed as a more robust definition of "opposed interests".

[*]Corresponding author.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

Our contributions in the context of related work. The central question we seek to address in this context is whether it is possible to "learn" in the presence of strictly opposed interests. Specifically, we seek to derive an adaptive, *uncoupled* learning scheme, where players update their actions without knowing the global structure of the game (or requiring payoff-related information of other players), and still converge to Nash equilibrium while achieving no regret.

There are several difficulties to overcome here. First, on the negative end of the spectrum, we show that there is no hope of leveraging an amenable "convex structure" in harmonic games. In stark contrast to min-max games, we show in Theorem 1 that, when the effective dimension of the game is odd, the set of fully mixed Nash equilibria is a *generically non-convex* real algebraic variety of dimension at least 1—a structural result of independent game-theoretic interest, established via the Poincaré–Hopf index theorem applied to a foliation of the game's strategy space by topological spheres. As such, standard convexity surrogates—like {quasi,star,quasar}-convexity [4, 17] or the related notions of variational stability [22, 27]—cannot be leveraged in this setting.

Second, in zero-sum games with a totally mixed equilibrium it is well known that standard no-regret learning dynamics—like *Follow-the-Regularized-Leader* (FTRL)—are Poincaré recurrent in continuous time [28, 35], and possibly divergent in discrete time [18, 29]. A standard fix to that is to employ an optimistic / extrapolation-based variant of FTRL, which "lenses" the trajectory of play and enables convergence to Nash equilibrium [16, 18, 29]—possibly even at a geometric rate, depending on the support of the equilibrium in question [46].

This picture was recently extended to harmonic games by Legacci et al. [24], who showed that, when run in a harmonic game, (*a*) FTRL is Poincaré recurrent in continuous time, (*b*) its vanilla implementation diverges (as in two-player zero-sum games), and (*c*) an extrapolation-based variant thereof converges asymptotically to Nash equilibrium, while (*d*) guaranteeing $\mathcal{O}(1)$ regret to all players.

Importantly, the learning scheme of Legacci et al. [24] is not quite uncoupled: while its iterative structure respects the uncoupledness constraints of Hart & Mas-Colell [19, 20], its parameter tuning does *not*. Each player must set their step-size with global knowledge of the game, including a specific equilibrium thereof, as well as information on the parameter configuration of all other players. In addition, the proposed regret guarantees in [24] are *brittle:* even if a single player deviates from the proposed algorithm, there is no guarantee that all other players will maintain even a worst-case $\mathcal{O}(\sqrt{T})$ regret.

Our starting point is precisely this: to show that none of these limitations are inherent. Specifically, we seek to design an *adaptive* learning algorithm that achieves convergence in harmonic games, all the while maintaining robust, "best-of-both-worlds" regret guarantees. We do this by introducing a player-specific version of the adaptive "difference of gradients squared" policy of Rakhlin & Sridharan [37, 38], which we combine with a general, variable-parameter extrapolated version of FTRL. In so doing, we show that (*i*) the generated trajectory of play converges to Nash equilibrium in all harmonic games; (*ii*) the proposed method achieves individual $\mathcal{O}(1)$ regret when all players follow the algorithm, implying in this way an $\mathcal{O}(1/T)$ rate of convergence of empirical play to coarse correlated equilibria; and (*iii*) the focal player is still guaranteed $\mathcal{O}(\sqrt{T})$ regret even against arbitrary play by all other players (including malicious interference of Byzantine faults). In this way, we simultaneously generalize the convergence results of Legacci et al. [24] to a fully adaptive setting, and extend the adaptive learning framework of [22, 38] to the class of harmonic games. [For further details on related work, cf. Appendix G.]

2 Basic definitions from game theory

Throughout the rest of our paper, we will focus on *finite games* consisting of (*i*) a finite number of *players* indexed by $i \in \mathcal{N} = \{1, \dots, N\}$; (*ii*) for each player $i \in \mathcal{N}$, a finite set of *actions*—or *pure strategies*—indexed by $\alpha_i \in \mathcal{A}_i = \{1, \dots, K_i\}$; and (*iii*) an ensemble of *payoff functions* $u_i : \mathcal{A} \to \mathbb{R}$, $i \in \mathcal{N}$, where $\mathcal{A} \coloneqq \prod_j \mathcal{A}_j$ denotes the set of all *action profiles* $\alpha = (\alpha_1, \dots, \alpha_N) \in \mathcal{A}_1 \times \cdots \times \mathcal{A}_N$. A finite game as above will be written in shorthand form as $\Gamma \equiv \Gamma(N, \mathcal{A}, u)$.

During play, each player $i \in \mathcal{N}$ may randomize their choice of action by means of a *mixed strategy* $x_i \in \Delta(\mathcal{A}_i)$, i.e., a probability distribution x_i over $\mathcal{A}_i$. We will write $x_{i\alpha_i}$ for the probability with which player $i \in \mathcal{N}$ plays action $\alpha_i \in \mathcal{A}_i$, $\mathcal{X}_i \coloneqq \Delta(\mathcal{A}_i)$ for the strategy space of player i, and $\mathcal{X} \coloneqq \prod_i \mathcal{X}_i$

for the game's strategy space, i.e., the space of all *mixed strategy profiles* $x = (x_1, \dots, x_N)$; also, when we want to focus on the strategy (or action) of a given player, we will use the shorthand $(x_i; x_{-i}) = (x_1, \dots, x_i, \dots, x_N)$, and likewise for pure strategies. In this manner, the *mixed payoff* to player $i \in \mathcal{N}$ in the strategy profile $x \in \mathcal{X}$ is given by

$$u_i(x) = \mathbb{E}_{\alpha \sim x}[u_i(\alpha)] = \sum_{\alpha_1 \in \mathcal{A}_1} \cdots \sum_{\alpha_N \in \mathcal{A}_N} x_{1,\alpha_1} \cdots x_{N,\alpha_N} u_i(\alpha_1, \dots, \alpha_N) \tag{1}$$

and we also write

$$v_i(x) = (u_i(\alpha_i; x_{-i}))_{\alpha_i \in \mathcal{A}_i} \quad \text{and} \quad v(x) = (v_1(x), \dots, v_N(x)) \quad \text{for all } x \in \mathcal{X} \tag{2}$$

for the *mixed payoff vector* of player $i \in \mathcal{N}$ and the associated *payoff field* of the game. By multilinearity, we have $u_i(x) = \langle v_i(x), x_i \rangle$ so $v(x)$ encodes all payoff-relevant information of the game; importantly, we also have $v_i(x) = \nabla_{x_i} u_i(x)$, so $v_i(x)$ can be seen as the direction of steepest individual payoff ascent.

For posterity, we recall that a strategy profile $x^* \in \mathcal{X}$ is a *Nash equilibrium* (NE) of Γ if

$$u_i(x^*) \geq u_i(x_i; x^*_{-i}) \qquad \text{for all } x_i \in \mathcal{X}_i, i \in \mathcal{N}. \tag{NE}$$

By definition, x^* is a Nash equilibrium if and only if it satisfies the (Stampacchia) variational inequality

$$\langle v(x^*), x - x^* \rangle \leq 0 \quad \text{for all } x \in \mathcal{X}. \tag{SVI}$$

In particular, if x^* is a pure strategy profile, we will refer to it as a *pure Nash equilibrium*; and if (NE) holds as a *strict* inequality for all $x_i \neq x^*_i$, $i \in \mathcal{N}$, we will refer to x^* as a *strict Nash equilibrium*. At the other end of the spectrum, if $x^*_{i\alpha_i} > 0$ for all $\alpha_i \in \mathcal{A}_i$, $i \in \mathcal{N}$—that is, if x^* has *full support*—we will say that x^* is *totally mixed*.

We detail below two classes of games that will play an important role in the sequel:

1. *Potential games:* A finite game $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$ is a potential game if it admits a potential function $\Phi \colon \mathcal{A} \to \mathbb{R}$ such that

   $$u_i(\alpha_i; \alpha_{-i}) - u_i(\beta_i; \alpha_{-i}) = \Phi(\alpha_i; \alpha_{-i}) - \Phi(\beta_i; \alpha_{-i}) \quad \text{for all } \alpha \in \mathcal{A}, i \in \mathcal{N}, \text{ and } \beta_i \in \mathcal{A}_i. \tag{3}$$

   Potential games represent strategic interactions with *common interests*, in the sense that the players' objectives are aligned along the game's potential Φ.

2. *Two-player zero-sum games:* As the name suggests, a zero-sum—or *min-max*—game is a finite game $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$ with $\mathcal{N} = \{1, 2\}$ (two players) and $u_1 = -u_2$ (zero-sum). Zero-sum games may have equilibria of any support, from pure to totally mixed; importantly, the set of Nash equilibria of a zero-sum game is always convex.

For obvious reasons, zero-sum games are the go-to model for adversarial strategic interactions: after all, one player's gain is the other's loss, and vice versa. Importantly however, the zero-sum property is *not* a strategic invariant of a game. The bimatrix game below is an example of a zero-sum game which is also a *potential* game; what's more, by shifting each player's payoffs by the same (irrelevant) amount when their opponent plays a specific strategy, the game is transformed into a cooperative, *common interest* one:

$$\begin{pmatrix} (0,0) & (1,-1) \\ (-1,1) & (0,0) \end{pmatrix} \sim \begin{pmatrix} (2^{(+2)},0) & (1,-1) \\ (1^{(+2)},1) & (0,0) \end{pmatrix} \sim \begin{pmatrix} (2,2^{(+2)}) & (1,1^{(+2)}) \\ (1,1) & (0,0) \end{pmatrix} \tag{4}$$

Instances like the above show that zero-sum games are not always faithful models of adversarial interactions and strictly opposed interests; we examine this issue in detail in the next section.

3 Harmonic games and their equilibria

In this section, we proceed to define the class of *harmonic games*, which will be our main model of adversarial strategic interactions. To begin, it will be instructive to provide a graphical description of a normal form game Γ by means of its *response graph*: this is a weighted directed graph $\mathrm{Resp}(\Gamma)$ with (*i*) vertices given by $\mathcal{A}$; (*ii*) a directed edge $\alpha \rightsquigarrow \beta$ whenever $\alpha_{-i} = \beta_{-i}$, $\alpha_i \neq \beta_i$, and $u_i(\beta) \geq u_i(\alpha)$ for some $i \in \mathcal{N}$ (that is, when the i-th player has a weak incentive to deviate from α_i to β_i); and (*iii*) a

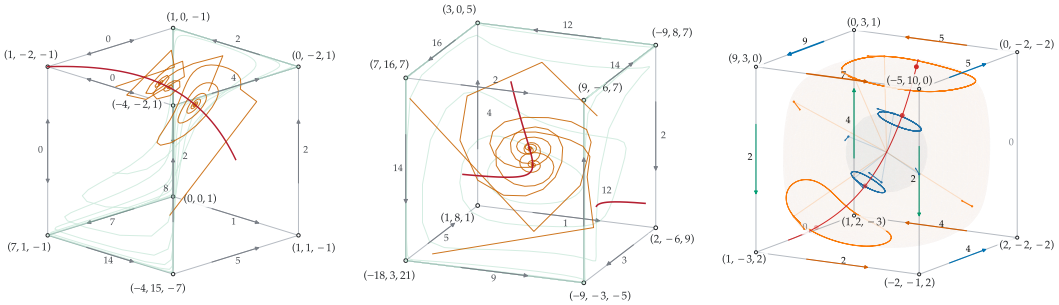

Figure 1: Response graphs, Nash equilibria (cf. Theorem 1), and dynamical outcomes (cf. Theorem 3) in three harmonic $2 \times 2 \times 2$ games. Red curves show 1-dimensional Nash sets. *Left:* a connected, non-convex Nash continuum. *Middle:* two disconnected, non-convex equilibrium branches. In both panels, light-green vanilla FTRL trajectories move outward, whereas orange adaptive FTRL+ trajectories converge to the Nash set. *Right:* blue and orange weighted Kullback–Leibler energy levels illustrate the foliation of $\mathrm{ri}\,\mathcal{X}$ by topological spheres, with tangent replicator trajectories and red markers at Nash intersections. The Poincaré–Hopf argument of Theorem 1 guarantees at least one equilibrium on every positive level.

weight $w_{\alpha\beta} \coloneqq u_i(\beta) - u_i(\alpha)$ assigned to each oriented edge $\alpha \rightsquigarrow \beta$. Clearly, the response graph of the game carries all payoff-relevant data of the game—in fact, two games Γ, Γ' are *strategically equivalent* if they have the same response graphs, that is, $\mathrm{Resp}(\Gamma) = \mathrm{Resp}(\Gamma')$.

An important property of potential games is that all directed paths in their response graphs terminate at a pure Nash equilibrium; in addition, every such game satisfies *Kirchhoff's voltage law*, i.e., the weight of any directed edge $\alpha \rightsquigarrow \beta$ is equal to the potential difference $\Phi(\beta) - \Phi(\alpha)$. Orthogonally to this, we say that a game is *harmonic* if it satisfies a (weighted version of) *Kirchhoff's current law*, that is, the total inflow in a vertex of $\mathcal{A}$ is equal to the total outflow from it, possibly up to a vertex-specific weight. Formally, we have the following definition.

Definition 1. A finite game $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$ is said to be *harmonic* when it admits a *harmonic measure*, i.e., a collection of weights $\mu_{i\alpha_i} \in (0, \infty)$, $\alpha_i \in \mathcal{A}_i$, $i \in \mathcal{N}$, such that

$$\sum_{\beta:\alpha\rightsquigarrow\beta} \mu_{\alpha\beta}\, w_{\alpha\beta} = \sum_{\beta:\beta\rightsquigarrow\alpha} \mu_{\alpha\beta}\, w_{\beta\alpha} \quad \text{for all } \alpha \in \mathcal{A}\,, \tag{HG}$$

where $\mu_{\alpha\beta} \coloneqq \mu_{i\beta_i}$ whenever $\alpha_i \neq \beta_i$ and $\alpha_{-i} = \beta_{-i}$. In particular, if Γ is harmonic relative to the uniform measure $\mu_{i\alpha_i} = 1$, $\alpha_i \in \mathcal{A}_i$, $i \in \mathcal{N}$, we will say that Γ is a *uniform harmonic game*.

The class of harmonic games was introduced by Candogan et al. [11]—and their definition extended by Abdou et al. [1]—who showed that, up to strategic equivalence, any finite game can be uniquely decomposed as the sum of a potential and a harmonic game (in the sense of addition of utilities). An immediate consequence of this result is that, in contrast to zero-sum games, the only game that can be simultaneously potential and harmonic up to strategic equivalence is the zero game.

Broadly speaking, in harmonic games, for any player considering a deviation toward a specific pure strategy profile, there exist other players with an incentive to deviate *away* from said profile. In conjunction with the decomposition result described above, harmonic games can thus be seen as "orthogonal" to potential games, and hence, as the canonical representatives of strategic conflict [1, 11]. Importantly, this does not contradict the long-standing use of games like Matching Pennies and Rock-Paper-Scissors as textbook examples of strictly opposed interests: these prototypical zero-sum games admit a totally mixed Nash equilibrium, and as we show in Appendix A, zero-sum games with a totally mixed NE *are* harmonic. In this regard, the adversarial nature of these classical games is more closely attuned to their harmonic character rather than the zero-sum property.

Beyond this case, zero-sum and harmonic games can differ sharply, even in their equilibrium structure. For pure equilibria, (4) shows that zero-sum games may admit a unique *strict* Nash equilibrium, whereas harmonic games admit none—see [11, Prop. 5.1] and [12, Thm. E.1].

For mixed equilibria, the picture is more subtle. In two-player zero-sum games, the minimax theorem guarantees that the Nash equilibrium set is convex. This amenable property extends to *two-player* harmonic games: Candogan et al. [11, Thm. B.2] showed that their set of mixed Nash equilibria is convex, with affine dimension $|K_1 - K_2|$. However, the geometry of this set beyond the two-player case remained an open problem, explicitly raised by Candogan et al. [11].

One property of the mixed-equilibrium set is known in full generality: Abdou et al. [1, Thm. 2.5] proved that every harmonic game admits a distinguished totally mixed Nash equilibrium $q \in \operatorname{ri} \mathcal{X}$, called the *harmonic center*; see Proposition A.1 for details.

Our first result makes substantial progress on this problem. Leveraging tools from differential topology, we show that the totally mixed Nash equilibria in odd-dimensional harmonic games form a real algebraic variety of dimension at least one, which need not be convex, and whose closure reaches the boundary of the strategy space:

Theorem 1 (Harmonic equilibria)**.** *Let $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$ be a harmonic game, and let $K = \sum_i K_i$ denote the total number of pure strategies in the game. If the game's effective dimension $d = K - N$ is odd, then the set of totally mixed Nash equilibria is a real algebraic variety of dimension at least* 1*, and its closure intersects the boundary* $\operatorname{bd} \mathcal{X}$ *of the game's strategy space $\mathcal{X}$.*

Sketch of proof of Theorem 1. The relative interior of any harmonic game's strategy space admits a one-parameter foliation by topological spheres centered at the harmonic center, to which the game's payoff vector field is tangent. In odd effective dimension, each leaf is then an even-dimensional sphere, on which a continuous tangent field must vanish—a consequence of the classical "hairy ball theorem" [3]. Each leaf therefore carries a zero of the payoff field, and hence a fully mixed Nash equilibrium; cf. Fig. 1–right. Boundary effects rule out a direct application of this argument: one then invokes the more general Poincaré–Hopf index theorem [3], applied to a suitable modification of the payoff field over a similar foliation. The complete proof is deferred to Appendix A.[2] ∎

Theorem 1 shows that the equilibrium structure of (odd-dimensional) harmonic games is quite exotic and highly non-generic. In addition, because such varieties are rarely convex (see Fig. 1 for two examples), this means that, in general—and in stark contrast to min-max [46] and variationally stable [22, 27] games—we cannot leverage convexity in order to design convergent learning algorithms in harmonic games. We explore the ramifications of this issue in the next sections.

4 The regularized learning template

Our learning framework builds on the seminal *Follow-the-Regularized-Leader* (FTRL) online learning template [42, 43]. We begin by reviewing some basic concepts in Section 4.1; subsequently, we present our learning template in Section 4.2.

4.1. Follow-the-Regularized-Leader. Let $\mathcal{V}$ be a normed vector space, $\mathcal{Y} \coloneqq \mathcal{V}^*$ its dual, and $\mathcal{X} \subseteq \mathcal{V}$ a closed convex set. In the online learning setting, an agent seeks to maximize their cumulative gain over $\mathcal{X}$ against a stream of linear payoffs $v_t \in \mathcal{Y}$, $t = 1, 2, \dots$, revealed sequentially. At each round, the Follow-the-Regularized-Leader algorithm plays the maximizer of the cumulative payoff up to time $t - 1$, regularized by a strongly convex penalty tailored to the geometry of $\mathcal{X}$. Concretely, the *Follow-the-Regularized-Leader* (FTRL) algorithm unfolds as

$$y_{t+1} = y_t + \gamma_t v_t \qquad x_{t+1} = Q(\eta_{t+1} y_{t+1}), \tag{FTRL$_{\gamma,\eta}$}$$

where the *mirror map* $Q \colon \mathcal{Y} \to \mathcal{X}$ is

$$Q(y) = \arg\max_{x \in \mathcal{X}} \{\langle y, x \rangle - h(x)\} \qquad \text{for all } y \in \mathcal{Y}, \tag{5}$$

and $h \colon \mathcal{X} \to \mathbb{R}$ is a strongly convex *regularizer* (cf. Definition 2 in Appendix B). The template (FTRL$_{\gamma,\eta}$) carries two distinct parameter sequences: a *step-size* $\gamma_t > 0$ that *pre*-multiplies each gradient as it enters the dual aggregate y_{t+1}, and a *learning rate* $\eta_t > 0$ that *post*-multiplies y_{t+1} before applying the mirror map. This dual schedule departs from the standard FTRL formulation—which features only a learning rate—but unifies several existing algorithms under a common umbrella, as detailed in Appendix C.1. We retain the name "FTRL" for the general template (FTRL$_{\gamma,\eta}$), with the caveat that different parameter schedules can lead to drastically different behaviors.

[2] An accompanying animation illustrates the proof: https://www.youtube.com/watch?v=do7mqQylEdc.

4.2. Regularized learning in games. To port the framework of the preceding section to our game-theoretic setting, it will be convenient to distinguish between (primal) "strategy-like" and (dual) "payoff-like" variables, writing $\mathcal{Y}_i \coloneqq \mathbb{R}^{K_i}$ and $\mathcal{Y} \coloneqq \prod_i \mathcal{Y}_i$ for the game's *payoff space*, in direct analogy with the game's strategy space $\mathcal{X} = \prod_i \mathcal{X}_i$. We assume further that each player is endowed with a κ_i-strongly convex regularizer $h_i \colon \mathcal{X}_i \to \mathbb{R}$ with induced mirror map $Q_i \colon \mathcal{Y}_i \to \mathcal{X}_i$; finally, for posterity, we will write $h(x) = \sum_i h_i(x_i)$ for the players' aggregate regularizer on $\mathcal{X}$, and $Q = (Q_1, \dots, Q_N) \colon \mathcal{Y} \to \mathcal{X}$ for the associated mirror map induced by h.

With all this in hand, building on the parent (FTRL$_{\gamma,\eta}$), we then obtain the following FTRL scheme:

$$y_{i,t+1} = y_{i,t} + \gamma_{i,t} v_i(x_t) \quad x_{i,t+1} = Q_i(\eta_{i,t+1} y_{i,t+1}) \quad \text{for each player } i \in \mathcal{N}. \tag{FTRL$_i$}$$

with variable step-size and learning rate parameters $\gamma_{i,t}, \eta_{i,t} > 0$. Unfortunately, a direct implementation of (FTRL$_i$) is known to fail in zero-sum games with totally mixed equilibria [8], ruling it out as a blueprint for harmonic games as well. Building instead on the extensive literature surrounding the extra-gradient algorithm of Korpelevich [23] and the "optimistic" variant of Popov [36], we will enrich FTRL with a "forward-looking" extrapolation step to help players anticipate their payoff landscape by emulating an implicit, proximal point update in the spirit of [26, 31, 40].

More precisely, we will consider the general *extrapolated FTRL* (FTRL+) scheme below, which updates in two distinct stages as follows:

$$\begin{aligned} &a)\ \textit{Extrapolate:} && y_{i,t+1/2} = y_{i,t} + \gamma_{i,t} \hat{v}_{i,t} && x_{i,t+1/2} = Q_i(\eta_{i,t+1/2} y_{i,t+1/2})\,. \\ &b)\ \textit{Update:} && y_{i,t+1} = y_{i,t} + \gamma_{i,t+1/2} \hat{v}_{i,t+1/2} && x_{i,t+1} = Q_i(\eta_{i,t+1} y_{i,t+1})\,. \end{aligned} \tag{FTRL+}$$

In the above, $\gamma_{i,t}, \eta_{i,t} > 0$, $t = 1/2, 1, \dots$, denote a non-increasing sequence of step-sizes and learning rates, and $\hat{v}_{i,t} \in \mathcal{Y}_i$ is a sequence of abstract "payoff signals", which we discuss below.

This is a very general template, so some remarks are in order. First, the method performs two FTRL updates: at half-integer times, an extrapolation step generates the *leading state* $x_{t+1/2}$ from x_t; at integer times, a second update from x_t produces the *base state* at time $t+1$, but using payoff information from $x_{t+1/2}$ rather than x_t—the algorithm's "anticipatory" component.

Second, as we mentioned earlier, the step-size and learning rate respectively pre- and post-multiply payoff information, playing vastly different roles: at a very high level, the former "lenses" the iterates near equilibrium, while the latter "stabilizes" them far from it. In some more detail, $\gamma_{i,t}$ determines how far the algorithm looks ahead before evaluating the payoff used for the actual update: near equilibrium, this look-ahead evaluation can detect the local rotational behavior of the payoff field, and introduce a corrective effect that can turn oscillation into contraction. Conversely, $\eta_{i,t}$ controls how sharply the accumulated score is translated into a strategy through the mirror map: when payoff predictions vary substantially, a small learning rate makes the iterates less sensitive to unstable or adversarial signals.

Last but not least is the role of the interleaved payoff signals $\hat{v}_t$ and $\hat{v}_{t+1/2}$. On this point, following Azizian et al. [7], a unifying feature of all extrapolation-based methods is that the information used to update the base state x_t must come *directly* from the method's leading state $x_{t+1/2}$, i.e.,

$$\hat{v}_{i,t+1/2} = v_i(x_{t+1/2}) \qquad \text{for all } i \in \mathcal{N} \text{ and all } t = 1, 2, \dots \tag{6a}$$

On the other hand, there is a plethora of choices for the extrapolation step itself—which we review for completeness in Appendix C.2. Mutatis mutandis, the extra-gradient [23], mirror-prox [6, 31], and dual extrapolation [32] methods all use $\hat{v}_{i,t} = v_i(x_t)$; by contrast, their optimistic variants [14, 21, 36, 38] use $\hat{v}_{i,t} = v_i(x_{t-1/2})$ instead. By this token, a flexible choice for $\hat{v}_t$ is the combination

$$\hat{v}_{i,t} = \lambda_i v_i(x_t) + (1 - \lambda_i) v_i(x_{t-1/2}) \qquad \text{for all } i \in \mathcal{N} \text{ and all } t = 1, 2, \dots \tag{6b}$$

for some ensemble of *extrapolation coefficients* $\lambda_i \in [0, 1]$, $i \in \mathcal{N}$.[3]

Importantly, all choices of γ_i, η_i and λ_i are player-specific, and we will not assume at any point in the sequel that they are tuned with information which is not available at the player side. With suitable choices of γ_t, η_t, and $\hat{v}_t$, (FTRL+) recovers a broad family of existing methods: as we show in Appendix C, *Mirror-prox* (MP) [6, 31], *Dual extrapolation* (DEX) [32], *Optimistic FTRL*

[3] We should note here that the "optimistic" update $\hat{v}_{i,t} \leftarrow v(x_{t-1/2})$ is more meaningful in online learning, as it relies exclusively on previous observations; in a game-theoretic setting however, both updates can be accommodated.

(OptFTRL) [2, 30, 44], and *Optimistic dual averaging* (OptDA) [22] all arise as special cases. Further specializations cover the conceptual mirror-prox of Nemirovski [31] with the implicit update $\hat{v}_t = v(x_{t+1})$, reflected-gradient methods [25], and, of course, vanilla FTRL (the case $\hat{v}_t = 0$). The closest precursor is Legacci et al. [24], whose version of (FTRL+) uses a single constant $\gamma_t = \gamma_{t+1/2} = \eta_t = \eta_{t+1/2} = c$—whence our notation and terminology. The decisive departure is our variable, multi-parameter schedule: this is what enables the unification above and unlocks adaptivity, at the price of a substantially more delicate analysis—to which we now turn.

5 Adaptive learning in harmonic games

We now turn to the analysis of (FTRL+). To streamline our presentation, we will assume in the rest of this section that the method is run with the default initialization scheme

$$y_1 \leftarrow 0\,, \quad x_1 \leftarrow x_c\,, \quad \eta_{i,1} \leftarrow 1\,, \quad \hat{v}_{i,1} \leftarrow \lambda_i v_i(x_c)\,, \tag{Init}$$

where $x_c \coloneqq Q(0) = \arg\min h$ denotes the so-called *prox-center* of $\mathcal{X}$ (usually its barycenter).

Groundwork. As we discuss in Appendix B, the *Fenchel coupling*

$$F(p, y) = h(p) + h^*(y) - \langle y, p\rangle \tag{7}$$

provides a primal-dual measure of distance between the primal points p and $Q(y)$. The cornerstone of our approach is a suitably rescaled—"η-*deflated*"—variant of this coupling, defined as

$$E_{i,t} = \frac{1}{\eta_{i,t}} F_i(p_i, \eta_{i,t} y_{i,t}) \qquad \text{for some } p_i \in \mathcal{X}_i, \text{ all } i \in \mathcal{N}, \text{ and all } t = 1/2, 1, \dots. \tag{8}$$

By Proposition B.2, $E_{i,t}$ provides a measure of distance between player i's strategy $x_{i,t} = Q_i(\eta_{i,t} y_{i,t})$ at time t and the base point $p_i \in \mathcal{X}_i$. The rescaling by $\eta_{i,t}$ "deflates" a spurious effect of the learning rate: a small $\eta_{i,t}$ pulls $x_{i,t}$ toward the prox-center $Q_i(0) = \arg\min h_i$ regardless of $y_{i,t}$, which would let the raw Fenchel coupling shrink without genuine progress; dividing by $\eta_{i,t}$ cancels this, so a decrease of $E_{i,t}$ unambiguously signals convergence toward p_i. Crucially, the deflated coupling satisfies a template inequality that lets us treat it as a quasi-Lyapunov function for (FTRL+):

Proposition 1 (Template inequality)**.** *Suppose that each player is running* (FTRL+) *with payoff signals as per* (6)*. Then, for all $p_i \in \mathcal{X}_i$, $i \in \mathcal{N}$, and all $t = 1, 2, \dots$, we have:*

$$E_{i,t+1} \leq E_{i,t} + \gamma_{i,t+1/2} \langle \hat{v}_{i,t+1/2}, x_{i,t+1/2} - p_i\rangle \tag{9a}$$

$$+ \left(\frac{1}{\eta_{i,t+1}} - \frac{1}{\eta_{i,t}}\right) \cdot [h_i(p_i) - \min h_i] \tag{9b}$$

$$+ \gamma_{i,t+1/2} \langle \hat{v}_{i,t+1/2} - \hat{v}_{i,t}, x_{i,t+1} - x_{i,t+1/2}\rangle \tag{9c}$$

$$- \frac{\gamma_{i,t+1/2}}{\gamma_{i,t}} \left[\frac{F_i(x_{i,t+1}, \eta_{i,t+1/2} y_{i,t+1/2})}{\eta_{i,t+1/2}} + \frac{F_i(x_{i,t+1/2}, \eta_{i,t+1/2} y_{i,t})}{\eta_{i,t+1/2}}\right] \tag{9d}$$

$$+ \frac{\gamma_{i,t+1/2}}{\gamma_{i,t}} \frac{F_i(x_{i,t+1}, \eta_{i,t+1/2} y_{i,t})}{\eta_{i,t+1/2}} - \frac{F_i(x_{i,t+1}, \eta_{i,t} y_{i,t})}{\eta_{i,t}}\,. \tag{9e}$$

Proposition 1 (proved in Appendix C.3) may appear fairly convoluted at first, but this is a consequence of the generality of (FTRL+)—a generality that surfaces a new "coupling" term (9e), emerging from the interplay between step-size and learning rate across the extrapolation and update phases. This coupling term can be ignored wholesale if we take

$$\gamma_{i,t+1/2} = \gamma_{i,t} = \gamma_i > 0 \qquad \text{and} \qquad \eta_{i,t+1/2} = \eta_{i,t}\,. \tag{10}$$

In fact, given the factor $\gamma_{i,t+1/2}/\gamma_{i,t}$ that appears in (9e), it might seem more meaningful to employ non-increasing step-size sequences fulfilling $\gamma_{i,t+1/2} \leq \gamma_{i,t}$, as this could yield a strictly negative contribution in (9). Nevertheless, a finer analysis [cf. the discussion following Eq. (D.2) in Appendix D] reveals that this additional control cannot be easily coupled with the other terms of (9). We therefore adopt the constant choice (10) to simplify both the analysis and the notation.

Convergence to equilibrium. A key consequence of (10) is that, after substituting in (9) and rearranging, the term (9a) readily becomes a regret proxy for (FTRL+). With this preliminary result in hand, we are in a position to state our first equilibrium convergence result for (FTRL+):

Theorem 2 (Convergence to equilibrium)**.** *Suppose that each player is running* (FTRL+) *in a harmonic game with payoff signals as per* (6) *and parameters as per* (10)*. If* $\liminf_t \eta_{i,t} > 0$ *and* $\gamma_i > 0$ *is sufficiently small, then* x_t *converges to the set of Nash equilibria of the game.*

Sketch of proof. The proof of Theorem 2 (which we detail in Appendix E) revolves around instantiating the template inequality (9) at the harmonic center q of the game. This has two consequences: First, the sequences $\|x_{t+1/2} - x_t\|^2$ and $\|x_{t+1} - x_{t+1/2}\|^2$ are both summable, so x_t and $x_{t+1/2}$ share the same limit points. Second, the total deflated Fenchel coupling $E_t \coloneqq \sum_i (m_i/\gamma_i) E_{i,t}$—weighted by the harmonic masses m_i divided by γ_i and instantiated at the harmonic center q—is quasi-Fejér [9], and hence converges to some finite limit E_∞. In turn, since $q \in \operatorname{ri} \mathcal{X}$, this implies that each scaled score $\eta_{i,t} y_{i,t}$ remains bounded modulo a constant shift in $\mathcal{Y}_i$, so, in turn, any limit point of x_t will lie in the prox-domain $\mathcal{X}_h$ of h. We then leverage the continuity of ∇h, the fact that $\liminf_t \eta_{i,t} > 0$ for all $i \in \mathcal{N}$, and the fact that x_t and x_{t+r} have the same limit points for any *fixed* $r \in \mathbb{N}$ to show that, if x_∞ is a limit point of x_t, the corresponding payoff vector $v(x_\infty)$ is outward pointing relative to $\mathcal{X}$ at x_∞, so x_∞ must be a Nash equilibrium. ∎

Stated thus, Theorem 2 extends a recent result by Legacci et al. [24], who provided a similar convergence guarantee for *constant* $\eta_{i,t}$. The extension to a variable learning rate requires handling several error terms in (9); however, this is not the main shortfall—nor the principal difficulty—of the theorem. The real obstruction lies in the meaning of "sufficiently small" for γ_i: as we detail in the proof of Proposition D.1 in Appendix D, the precise threshold is

$$\gamma_i \leq \min\{1, \kappa_i \cdot (\min_j m_j \min_j \kappa_j) \big/ (16\, m_i N L_i^2)\}\,, \tag{11}$$

which depends on quantities requiring global knowledge of the game (the harmonic masses m_i, the Lipschitz moduli L_i, the strong-convexity moduli κ_i, and the player count N) that no individual player can observe. Compounding this, Theorem 2 also requires $\liminf_t \eta_{i,t} > 0$, ruling out the vanishing schedules that are typical of regret minimization. In short, Theorem 2—as well as the corresponding result of Legacci et al. [24]—is of limited applicability, as it requires each player to commit, ex ante, to constants chosen from full prior knowledge of the game.

To account for this, we will instead consider an adaptive "difference of gradients squared" learning rate policy in the spirit of Rakhlin & Sridharan [38], and run (FTRL+) with

$$\eta_{i,t} = \frac{1}{\sqrt{1 + \sum_{s=1}^{t-1} \|\hat{v}_{i,s+1/2} - \hat{v}_{i,s}\|_*^2}}\,. \tag{12}$$

The intuition behind (12) is the ansatz that, if x_t converges to the set of Nash equilibria of the game, the payoff signals $\hat{v}_t$ and $\hat{v}_{t+1/2}$ will equilibrate to the same value; as such, the series in (12) will plausibly converge to some finite limit as $t \to \infty$. Moreover, since the learning rate is now defined in terms of successive payoff differences, it automatically tracks the variation of observed payoff data, so it can be expected to adjust to the local Lipschitz modulus of v_i—which appears in the bound (11). By contrast, if the iterates of (FTRL+) are far from Nash equilibrium, $\eta_{i,t}$ will decrease (at least for a period of time) as $\Theta(1/\sqrt{t})$, the canonical rate for regret-based analyses. Finally, one last point worth emphasizing is that the schedule (12) is entirely environment-agnostic, relying only on information *locally* available to each player.

In this adaptive context, we have the following equilibrium convergence guarantee for (FTRL+).

Theorem 3 (Adaptive convergence)**.** *Suppose that each player is running* (FTRL+) *in a harmonic game with payoff signals as per* (6)*, parameters as per* (10) *and* (12)*, and* $\gamma_i = 1$*. Then* x_t *converges to the set of Nash equilibria of the game.*

Sketch of proof. The key difficulty in the proof of Theorem 3 is to show that $\lim_{t\to\infty} \eta_{i,t} > 0$ for all $i \in \mathcal{N}$; we do this via a summability analysis combined with a stopping time argument. As detailed in Proposition D.1 and Appendix E, this again revolves around summing the template inequality (9) instantiated at the game's harmonic center q and weighted by the game's harmonic masses m_i. However, this point is not enough to invoke Theorem 2, because the "small steps" condition (11) is

not satisfied. We thus leverage the adaptive structure of (12) to show independently that the sequences $\|x_{t+1/2} - x_t\|^2$ and $\|x_{t+1} - x_{t+1/2}\|^2$ are both summable, and that the total deflated Fenchel coupling $E_t = \sum_i m_i E_{i,t}$ instantiated at the harmonic center q is quasi-Fejér. Armed with these results, we can tie back into the proof blueprint of Theorem 2, which ultimately yields the desired claim. ∎

The "parameter-agnostic" requirement of (12) is very natural from an operational standpoint, yet only a handful of works in the learning-in-games literature satisfy it. The seminal work of Rakhlin & Sridharan [38] employs the same policy to obtain *average* convergence to minimax equilibria in zero-sum games—a strictly weaker guarantee than the last-iterate convergence of Theorem 3. The closest comparable results to Theorem 3 are those of Hsieh et al. [22], who study (OptDA) in *variationally stable* games with an adaptive learning rate; and of Legacci et al. [24], for (FTRL+) with *constant* η_i in harmonic games. Our work advances on both, structurally and technically.

First and foremost, despite a formal resemblance between Eq. (A.1) and variational stability (VS), harmonic games are *not* variationally stable: VS implies a convex set of Nash equilibria [27, Prop. 2.7], whereas harmonic games exhibit genuinely non-convex equilibrium varieties. For a concrete example, Hsieh et al. [22]'s analysis fails on the 3-player harmonic games of Fig. 1. Theorem 1 thus places harmonic games in a class strictly harder than variationally stable or zero-sum games—one in which no convexity surrogate is available, and where different proof techniques are required. The proof of Theorem 3 (Appendix E) accordingly departs from [22], anchoring the convergence analysis at the game's harmonic center. Theorem 3 also lifts a key assumption in [22, Thm. 7], which guarantees Nash convergence only under *strict* VS (a consistent drift *toward* equilibrium) or under VS with a *non-steep* regularizer. The latter excludes the central setting of zero-sum games with fully mixed equilibrium under entropic regularization, where Hsieh et al. [22, Thm. 8] resort to a dedicated, ad-hoc argument; our analysis subsumes that result and extends it to N-player harmonic games and arbitrary regularizers, steep or not.

A further technical hurdle, absent from [22, 24, 38], is that the extrapolation phase of (FTRL+) accommodates convex combinations of extra-gradient and optimistic updates, requiring fine control of the gradient differences in the adaptive policy (12)—which we achieve by means of Lemma D.1.

Regret guarantees. The last question we treat is one of robustness, in two senses: individual payoff guarantees, and Byzantine fault tolerance. Recall that the regret of player $i \in \mathcal{N}$ against a sequence $X_t \in \mathcal{X}$, $t \in \mathbb{N}$, is

$$\mathrm{Reg}_i(T) = \max_{p_i \in \mathcal{X}_i} \sum_{t=1}^{T} [u_i(p_i; X_{-i,t}) - u_i(X_t)] \,. \tag{13}$$

We then obtain the following guarantee for the sequence $X_{i,t} = x_{i,t+1/2}$ generated by (FTRL+), in two regimes: (a) the focal player runs the algorithm in an arbitrary environment; and (b) all players do, and the game is harmonic.

Theorem 4 (Individual regret guarantees)**.** *Suppose that player $i \in \mathcal{N}$ is running* (FTRL+) *with payoff signals as per* (6)*, parameters as per* (10) *and* (12)*, and $\gamma_i = 1$. Then, letting $H_i \coloneqq \max h_i - \min h_i$ and $G_i = \sup_{x \in \mathcal{X}} \|v_i(x)\|_*$, we have:*

(a) The regret of player $i \in \mathcal{N}$ against any sequence $X_{-i,t} \in \mathcal{X}_{-i}$, $t = 1, 2, \dots$, is bounded as

$$\mathrm{Reg}_i(T) \leq (H_i + 1/\kappa_i)(1 + 4G_i^2 T)^{1/2} + 2\,G_i^2 \,/\, \kappa_i = \mathcal{O}\big(\sqrt{T}\big) \,. \tag{14a}$$

(b) If, in addition, all players follow (FTRL+) *as above and the game is harmonic, then*

$$\mathrm{Reg}_i(T) = \mathcal{O}(H_i + 1/\kappa_i) \,. \tag{14b}$$

The proof of Theorem 4 is deferred to Appendix D. We note here that the bound (14a) matches the rate of the analogous result in [22] while applying to a more general algorithmic template. The bound (14b) is new: the closest analogue in [24] relies on a *constant* schedule that (a) must be tuned from full knowledge of the game; and (b) exposes the focal player to *linear* regret against arbitrary play. We illustrate Theorem 4 and the brittleness it lifts relative to Legacci et al. [24] numerically in Section 6.

6 Numerical experiments

We test adaptive FTRL+ on the two three-player harmonic games in Fig. 1 (left and middle), using entropic regularization and $\lambda_i = 1/2$. Across ten interior initializations per game, the median Nash gap decays rapidly despite the non-convexity and, in the second example, disconnectedness of the equilibrium set (Fig. 2a–b). The interquartile bands show variation across initializations. The approximately straight decline of the log Nash gap against rounds is compatible with linear (geometric) convergence of adaptive FTRL+ in the two harmonic examples of Fig. 1.

Unilateral robustness of Theorem 4 is illustrated on the disconnected example by comparing constant and adaptive policies, both in self-play and with a third player using an oversized constant score step. In self-play, both regret histories remain nearly flat. With a deviator, constant-policy regret grows approximately linearly, whereas adaptive regret remains nearly flat over 8,000 rounds (Fig. 2c), illustrating the robustness guarantee of Theorem 4-(a). Parameters and initialization conventions are given in Appendix F.

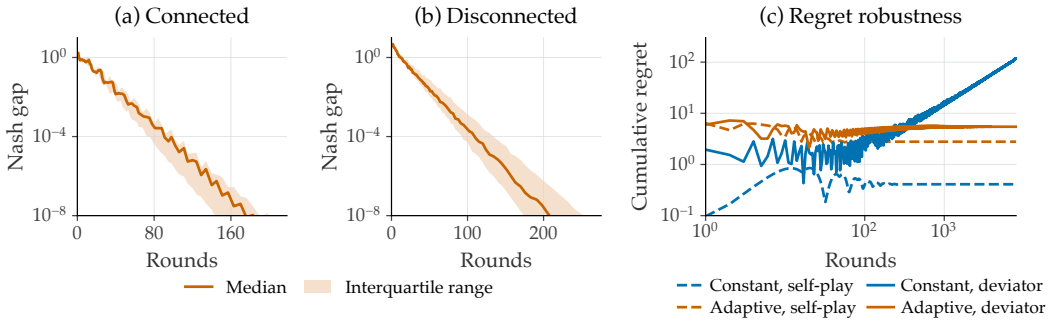

Figure 2: Convergence rates and the robustness trade-off of FTRL+. *(a,b)* Log Nash gap over rounds over ten initializations in the harmonic games of Fig. 1. *(c)* Both the constant and adaptive policies exhibit essentially constant cumulative regret in self-play, with a better constant for the constant-step policy. Against the same deviating policy and initialization, however, the maximum regret of constant-step players grows approximately linearly, while adaptive regret remains nearly constant. Adaptivity thus trades a larger self-play regret constant for robustness guarantee against nonconforming play.

7 Concluding remarks

The two existing strands of analysis leave a structural gap: adaptive methods such as [22] dispense with parameter tuning, but do not cover harmonic games; methods convergent in harmonic games [24] require global problem parameters—including the harmonic weights—and offer no worst-case regret guarantees. Closing this gap matters: harmonic games are a robust model of N-player strategic conflict beyond the zero-sum paradigm; adaptive, regret-optimal dynamics in this class are therefore an important step toward scaling and stabilizing competitive multi-agent learning beyond the stable and saddle-point regimes. (FTRL+) delivers exactly this with a flexible, parameter-agnostic template: last-iterate convergence to Nash equilibrium, $\mathcal{O}(1)$ regret in self-play, and $\mathcal{O}(\sqrt{T})$ regret against arbitrary play. This answers in the positive an open question posed by Legacci et al. [24], and shows that (FTRL+) is particularly appealing as an adaptive learning template for games with conflicting strategic interests.

Two questions remain open. First, our analysis assumes full-information feedback; whether comparable guarantees extend to bandit or payoff-only feedback remains unknown. Second, while our convergence guarantee is parameter-agnostic, it is setwise and provides no rate.

Numerical evidence in Fig. 2 suggests that FTRL+ may achieve linear convergence for general harmonic games, and without tuning. However, establishing rates appears genuinely hard: harmonic games lack the monotonicity, operator linearity, and equilibrium convexity that underpin standard error-bound arguments [46]. Obtaining adaptive learning with guaranteed linear convergence in all harmonic games remains a promising direction for future research.

Acknowledgments and Disclosure of Funding

This research was supported in part by the French National Research Agency (ANR) in the framework of the PEPR IA FOUNDRY project (ANR-23-PEIA-0003). PM is also a member of the Archimedes Research Unit/Athena RC, and was partially supported by project MIS 5154714 of the National Recovery and Resilience Plan Greece 2.0 funded by the European Union under the NextGenerationEU Program.

References

[1] Abdou, J., Pnevmatikos, N., Scarsini, M., and Venel, X. Decomposition of games: Some strategic considerations. *Mathematics of Operations Research*, 47(1):176–208, February 2022.

[2] Abernethy, J., Lai, K. A., Levy, K. Y., and Wang, J.-K. Faster rates for convex-concave games. In *COLT '18: Proceedings of the 31st Annual Conference on Learning Theory*, 2018.

[3] Armstrong, M. A. *Basic Topology*. Springer, 1983.

[4] Arrow, K. J. and Enthoven, A. C. Quasi-concave programming. *Econometrica*, 29(4):779–800, 1961.

[5] Auer, P., Cesa-Bianchi, N., and Gentile, C. Adaptive and Self-Confident On-Line Learning Algorithms. *Journal of Computer and System Sciences*, 64(1):48–75, 2002.

[6] Auslender, A. and Teboulle, M. Interior projection-like methods for monotone variational inequalities. *Mathematical Programming*, 104:39–68, 2005.

[7] Azizian, W., Iutzeler, F., Malick, J., and Mertikopoulos, P. The rate of convergence of Bregman proximal methods: Local geometry vs. regularity vs. sharpness. *SIAM Journal on Optimization*, 34(3):2440–2471, September 2024.

[8] Bailey, J. P. and Piliouras, G. Multiplicative Weights Update in Zero-Sum Games. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, pp. 321–338, New York, NY, USA, June 2018. Association for Computing Machinery.

[9] Bauschke, H. H. and Combettes, P. L. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, New York, NY, USA, 2 edition, 2017.

[10] Bertsekas, D. P. *Convex optimization algorithms*. Athena Scientific, 2015.

[11] Candogan, O., Menache, I., Ozdaglar, A., and Parrilo, P. A. Flows and decompositions of games: Harmonic and potential games. *Mathematics of Operations Research*, 36(3):474–503, 2011.

[12] Cauvin, P.-L., Legacci, D., and Mertikopoulos, P. The impact of uncertainty on regularized learning in games. In *ICML '25: Proceedings of the 42nd International Conference on Machine Learning*, 2025.

[13] Chen, G. and Teboulle, M. Convergence analysis of a proximal-like minimization algorithm using Bregman functions. *SIAM Journal on Optimization*, 3(3):538–543, August 1993.

[14] Chiang, C.-K., Yang, T., Lee, C.-J., Mahdavi, M., Lu, C.-J., Jin, R., and Zhu, S. Online optimization with gradual variations. In *COLT '12: Proceedings of the 25th Annual Conference on Learning Theory*, 2012.

[15] Combettes, P. L. Quasi-Fejérian analysis of some optimization algorithms. In Butnariu, D., Censor, Y., and Reich, S. (eds.), *Inherently Parallel Algorithms in Feasibility and Optimization and Their Applications*, pp. 115–152. Elsevier, New York, NY, USA, 2001.

[16] Daskalakis, C. and Panageas, I. Last-iterate convergence: Zero-sum games and constrained min-max optimization. In *ITCS '19: Proceedings of the 10th Conference on Innovations in Theoretical Computer Science*, 2019.

[17] de Finetti, B. Sulle stratificazioni convesse. *Annali di Matematica Pura ed Applicata*, 30:173–183, 1949.

[18] Gidel, G., Berard, H., Vignoud, G., Vincent, P., and Lacoste-Julien, S. A variational inequality perspective on generative adversarial networks. In *ICLR '19: Proceedings of the 2019 International Conference on Learning Representations*, 2019.

[19] Hart, S. and Mas-Colell, A. Uncoupled dynamics do not lead to Nash equilibrium. *American Economic Review*, 93(5): 1830–1836, 2003.

[20] Hart, S. and Mas-Colell, A. Stochastic uncoupled dynamics and Nash equilibrium. *Games and Economic Behavior*, 57: 286–303, 2006.

[21] Hsieh, Y.-G., Iutzeler, F., Malick, J., and Mertikopoulos, P. On the convergence of single-call stochastic extra-gradient methods. In *NeurIPS '19: Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pp. 6936–6946, 2019.

[22] Hsieh, Y.-G., Antonakopoulos, K., and Mertikopoulos, P. Adaptive learning in continuous games: Optimal regret bounds and convergence to Nash equilibrium. In *COLT '21: Proceedings of the 34th Annual Conference on Learning Theory*, 2021.

[23] Korpelevich, G. M. The extragradient method for finding saddle points and other problems. *Èkonom. i Mat. Metody*, 12: 747–756, 1976.

[24] Legacci, D., Mertikopoulos, P., Papadimitriou, C. H., Piliouras, G., and Pradelski, B. S. R. No-regret learning in harmonic games: Extrapolation in the presence of conflicting interests. In *NeurIPS '24: Proceedings of the 38th International Conference on Neural Information Processing Systems*, 2024.

[25] Malitsky, Y. Projected reflected gradient methods for monotone variational inequalities. *SIAM Journal on Optimization*, 25(1):502–520, 2015.

[26] Martinet, B. Régularisation d'inéquations variationnelles par approximations successives. *ESAIM: Mathematical Modelling and Numerical Analysis*, 4(R3):154–158, 1970.

[27] Mertikopoulos, P. and Zhou, Z. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, 173(1-2):465–507, January 2019.

[28] Mertikopoulos, P., Papadimitriou, C. H., and Piliouras, G. Cycles in adversarial regularized learning. In *SODA '18: Proceedings of the 29th annual ACM-SIAM Symposium on Discrete Algorithms*, 2018.

[29] Mertikopoulos, P., Lecouat, B., Zenati, H., Foo, C.-S., Chandrasekhar, V., and Piliouras, G. Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile. In *ICLR '19: Proceedings of the 2019 International Conference on Learning Representations*, 2019.

[30] Mohri, M. and Yang, S. Accelerating online convex optimization via adaptive prediction. In *AISTATS '16: Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, 2016.

[31] Nemirovski, A. S. Prox-method with rate of convergence $O(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.

[32] Nesterov, Y. Dual extrapolation and its applications to solving variational inequalities and related problems. *Mathematical Programming*, 109(2):319–344, 2007.

[33] Nesterov, Y. Primal-dual subgradient methods for convex problems. *Mathematical Programming*, 120(1):221–259, 2009.

[34] Orabona, F. and Pál, D. Scale-free online learning. *Theoretical Computer Science*, 716:50–69, 2018.

[35] Piliouras, G. and Shamma, J. S. Optimization despite chaos: Convex relaxations to complex limit sets via Poincaré recurrence. In *SODA '14: Proceedings of the 25th annual ACM-SIAM Symposium on Discrete Algorithms*, 2014.

[36] Popov, L. D. A modification of the Arrow–Hurwicz method for search of saddle points. *Mathematical Notes of the Academy of Sciences of the USSR*, 28(5):845–848, 1980.

[37] Rakhlin, A. and Sridharan, K. Online learning with predictable sequences. In *COLT '13: Proceedings of the 26th Annual Conference on Learning Theory*, 2013.

[38] Rakhlin, A. and Sridharan, K. Optimization, learning, and games with predictable sequences. In *NIPS '13: Proceedings of the 27th International Conference on Neural Information Processing Systems*, 2013.

[39] Rockafellar, R. T. *Convex Analysis*. Princeton University Press, Princeton, NJ, 1970.

[40] Rockafellar, R. T. Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14(5):877–898, 1976.

[41] Rockafellar, R. T. and Wets, R. J. B. *Variational Analysis*, volume 317 of *A Series of Comprehensive Studies in Mathematics*. Springer-Verlag, Berlin, 1998.

[42] Shalev-Shwartz, S. *Online learning: Theory, algorithms, and applications*. PhD thesis, Hebrew University of Jerusalem, 2007.

[43] Shalev-Shwartz, S. and Singer, Y. Convex repeated games and Fenchel duality. In *NIPS' 06: Proceedings of the 19th Annual Conference on Neural Information Processing Systems*, pp. 1265–1272. MIT Press, 2006.

[44] Syrgkanis, V., Agarwal, A., Luo, H., and Schapire, R. E. Fast convergence of regularized learning in games. In *NIPS '15: Proceedings of the 29th International Conference on Neural Information Processing Systems*, pp. 2989–2997, 2015.

[45] von Neumann, J. and Morgenstern, O. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

[46] Wei, C.-Y., Lee, C.-W., Zhang, M., and Luo, H. Linear last-iterate convergence in constrained saddle-point optimization. In *ICLR '21: Proceedings of the 2021 International Conference on Learning Representations*, 2021.

[47] Xiao, L. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research*, 11:2543–2596, October 2010.

Appendices

A Harmonic games and their equilibria **13**

A.1 A few facts about harmonic games 13

A.2 Proof of Theorem 1 14

B Properties of regularizers and mirror maps **15**

B.1 Annihilator and normal cone 15

B.2 Bregman regularization 15

C Algorithm analysis **18**

C.1 Mirror descent and its friends 18

C.2 Relation between (FTRL+) and other extrapolation methods 19

C.3 Template inequality for (FTRL+) 20

D Regret analysis **21**

D.1 Adversarial regret analysis 21

D.2 Analysis of regret in harmonic games 23

E Convergence analysis **27**

E.1 Preliminary results 27

E.2 Quotient dual space 29

E.3 Convergence to the set of Nash equilibria 29

F Additional numerical details **31**

G Further related work **32**

G.1 Extrapolation in games 32

G.2 Adaptive methods in games 32

G.3 Convergence in harmonic games 32

A Harmonic games and their equilibria

A.1. A few facts about harmonic games. Legacci et al. [24, Prop. A.4] first showed the following characterization of harmonic games in terms of mixed strategies:

Proposition A.1 (Harmonic center)**.** *A finite game $\Gamma(\mathcal{N}, \mathcal{A}, u)$ is harmonic with measure μ if and only if there exist a totally mixed strategy $q \in \mathcal{X}$ and, for each player, a* mass $m_i = \sum_{\alpha_i \in \mathcal{A}_i} \mu_{i\alpha_i} > 0$ *such that*

$$\sum_{i \in \mathcal{N}} m_i \, \langle v_i(x), x_i - q_i \rangle = 0 \quad \text{for all } x \in \mathcal{X} \,. \tag{A.1}$$

Proof. It is straightforward to verify that a game is harmonic if and only if $F(\alpha) \coloneqq \sum_{i \in \mathcal{N}} F_i(\alpha) = 0$ for all $\alpha \in \mathcal{A}$, where $F_i(\alpha) \coloneqq \sum_{\beta_i \in \mathcal{A}_i} \mu_{i\beta_i} [u_i(\alpha_i; \alpha_{-i}) - u_i(\beta_i; \alpha_{-i})]$. By multilinearity, this is equivalent to the condition $\sum_{i \in \mathcal{N}} m_i \langle v_i(x), x_i - \mu_i/m_i \rangle = 0$. ∎

The totally mixed strategy q appearing in Eq. (A.1) is called *harmonic center* of the game. As first shown by Abdou et al. [1, Thm. 1], the harmonic center is always a Nash equilibrium:

Proposition A.2. *Every harmonic game admits its harmonic center q as a totally mixed Nash equilibrium.*

Proof. Let q be the harmonic center of the game. Fix any $j \in \mathcal{N}$ and $x_j \in \mathcal{X}_j$, and let $x = (x_j, q_{-j})$. Since by multilinearity the mixed payoff vector of any player does not depend on the strategies of said player, we have that $v_j(x) = v_j(x_j, q_{-j}) = v_j(q)$. Instantiating Eq. (A.1) at x then gives

$$0 = m_j \langle v_j(x), x_j - q_j \rangle + \sum_{i \neq j} m_i \langle v_i(x), x_i - q_i \rangle = m_j \langle v_j(q), x_j - q_j \rangle . \tag{A.2}$$

Thus, for all $j \in \mathcal{N}$ and $x_j \in \mathcal{X}_j$, (NE) holds with equality, that is, $u_j(q) = u_j(x_j, q_{-j})$. ∎

Moving on, Legacci et al. [24, Prop. A.10] showed that every two-player zero-sum game with a totally mixed Nash equilibrium x^* is harmonic, with harmonic measure $\mu = x^*$:

Proposition A.3. *Every two-player zero-sum game with a totally mixed Nash equilibrium x^* is harmonic, with measure $\mu_{i\alpha_i} = x^*_{i\alpha_i}$ for each $i \in \mathcal{N}$ and $\alpha_i \in \mathcal{A}_i$.*

Proof. Let Γ be a two-player zero-sum game. For any $x \in \mathcal{X}$ it follows by bilinearity that

$$\begin{aligned} \sum_{i \in \mathcal{N}} \langle v_i(x), x_i - x_i^* \rangle &= \langle v_1(x), x_1 - x_1^* \rangle + \langle v_2(x), x_2 - x_2^* \rangle \\ &= u_1(x_1, x_2) - u_1(x_1^*, x_2) + u_2(x_1, x_2) - u_2(x_1, x_2^*) = 0 , \end{aligned}$$

where in the second line we used the assumption that x^* is a totally mixed Nash equilibrium; the conclusion then follows by Proposition A.1. ∎

A.2. Proof of Theorem 1. We now turn to the proof of Theorem 1, which we restate below for convenience.

Theorem 1 (Harmonic equilibria)**.** *Let $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$ be a harmonic game, and let $K = \sum_i K_i$ denote the total number of pure strategies in the game. If the game's effective dimension $d = K - N$ is odd, then the set of totally mixed Nash equilibria is a real algebraic variety of dimension at least 1, and its closure intersects the boundary $\operatorname{bd} \mathcal{X}$ of the game's strategy space $\mathcal{X}$.*

Proof. Let q denote the harmonic center of the game, so $\sum_{i \in \mathcal{N}} m_i \langle v_i(x), x_i - q_i \rangle = 0$ for all $x \in \mathcal{X}$ by Proposition A.1. Consider also the "μ-weighted" Kullback–Leibler divergence relative to q

$$E(q, x) = \sum_{i \in \mathcal{N}} m_i \sum_{\alpha_i \in \mathcal{A}_i} q_{i\alpha_i} \log \frac{q_{i\alpha_i}}{x_{i\alpha_i}} \quad \text{for all } x \in \operatorname{ri} \mathcal{X} , \tag{KL}$$

and the associated replicator field $\mathrm{RF} \colon \mathcal{X} \to \mathbb{R}^K$ induced by v, viz.

$$\mathrm{RF}_{i\alpha_i}(x) = x_{i\alpha_i} [v_{i\alpha_i}(x) - u_i(x)] . \tag{RF}$$

By a straightforward differentiation, we readily get

$$\frac{\partial \, \mathrm{KL}}{\partial x_{i\alpha_i}} = -m_i \frac{q_{i\alpha_i}}{x_{i\alpha_i}} \tag{A.3}$$

and thus, using that $\sum_{\alpha_i} q_{i\alpha_i} = 1$ and $u_i(x) = \langle v_i(x), x_i \rangle$ for all $i \in \mathcal{N}$ and $\alpha_i \in \mathcal{A}_i$,

$$\begin{aligned} \langle \nabla \, \mathrm{KL}(x), \mathrm{RF}(x) \rangle &= - \sum_{i \in \mathcal{N}} m_i \sum_{\alpha_i \in \mathcal{A}_i} \frac{q_{i\alpha_i}}{x_{i\alpha_i}} x_{i\alpha_i} [v_{i\alpha_i}(x) - u_i(x)] \\ &= \sum_{i \in \mathcal{N}} m_i \langle v_i(x), x_i - q_i \rangle = 0 \end{aligned} \tag{A.4}$$

for all $x \in \mathcal{X}$. This shows that the replicator field $\mathrm{RF}(x)$ is tangent to the level sets

$$\mathcal{S}_r = \{ x \in \mathcal{X} : E(q, x) = r \} \tag{A.5}$$

of $E(q, x)$. Now, since E is convex in x (as the sum of convex functions), it follows that the corresponding sublevel sets $\mathcal{S}_r^-(r) = \{ x \in \mathcal{X} : E(q, x) \leq r \}$ are convex and compact, which in turn

implies that each $\mathcal{S}_r$ is a topological sphere (the boundary of a smooth convex set), for all $r > 0$. Accordingly, since $\mathrm{RF}(x)$ is tangent to $\mathcal{S}_r$ whenever $x \in \mathcal{S}_r$, we conclude by the "hairy ball" version of the Poincaré–Hopf index theorem that RF admits a zero point $x_r^* \in \mathcal{S}_r$ for all $r > 0$. Since the level sets $\mathcal{S}_r$ are disjoint for different values of r, and the interior zeros of $\mathrm{RF}(x)$ are, a fortiori, Nash equilibria of the game, it follows that the game's set of totally mixed equilibria contains a 1-parameter subset. Since the set of totally mixed equilibria is a real algebraic variety and $\operatorname{ri}\mathcal{X} = \bigcup_{r\geq 0}\mathcal{S}_r$, our assertion follows. ∎

B Properties of regularizers and mirror maps

In this appendix we collect several standard properties concerning regularizers and the induced maps. In the following $\mathcal{V}$ denotes a normed vector space with norm $\|\cdot\|$, $\mathcal{V}^*$ denotes its dual, and $\mathcal{C} \subseteq \mathcal{V}$ is assumed closed and convex. The properties relevant to the game-theoretic setting follow by specializing $\mathcal{V}$ to $\mathbb{R}^{\mathcal{A}_i}$ or $\prod_{j\in\mathcal{N}} \mathbb{R}^{\mathcal{A}_j}$, and $\mathcal{C}$ to a player's strategy space $\mathcal{X}_i$ or to the product thereof.

B.1. Annihilator and normal cone. Let $\operatorname{aff}\mathcal{C}$ be the affine hull of $\mathcal{C}$, $T \subseteq \mathcal{V}$ the vector subspace of $\mathcal{V}$ parallel to $\operatorname{aff}\mathcal{C}$, and $T^\circ \subseteq \mathcal{V}^*$ the annihilator of T:

$$T = \operatorname{aff}\mathcal{C} - \operatorname{aff}\mathcal{C} = \{p - x : p, x \in \mathcal{C}\} \subseteq \mathcal{V}\,, \tag{B.1a}$$

$$T^\circ = \{v \in \mathcal{V}^* : \langle v, p - x\rangle = 0 \ \text{ for all } p, x \in \mathcal{C}\}\,. \tag{B.1b}$$

For any $x \in \mathcal{C}$, the *normal cone* to $\mathcal{C}$ is

$$\mathrm{NC}_x = \{v \in \mathcal{V}^* : \langle v, p - x\rangle \leq 0 \ \text{ for all } p \in \mathcal{C}\}\,. \tag{B.2}$$

Clearly, $T^\circ \subseteq \mathrm{NC}_x$ for all $x \in \mathcal{C}$. If x is in the relative interior of $\mathcal{C}$, the opposite is also true:

Lemma B.1. *If $x \in \operatorname{ri}\mathcal{C}$ then $T^\circ \supseteq \mathrm{NC}_x$. In particular, $T^\circ = \mathrm{NC}_x$.*

Proof. If $x \in \operatorname{ri}\mathcal{C}$ then every tangent vector is locally feasible in both directions: for all $d \in T$ there exist $\epsilon > 0$ such that $x + td \in \mathcal{C}$ for all $|t| \leq \epsilon$. If $v \in \mathrm{NC}_x$, then $\langle v, x \pm \epsilon d - x\rangle \leq 0$, so $v \in T^\circ$. ∎

B.2. Bregman regularization. Let $h\colon \mathcal{V} \to \mathbb{R} \cup \{\infty\}$ be a convex function, finite and continuous on $\mathcal{C}$, and denote by $h|_{\mathcal{C}}$ its restriction to $\mathcal{C}$. By assumption, $\mathcal{C} \subseteq \operatorname{dom} h = \{x \in \mathcal{V} : h(x) < \infty\}$; hence, $\mathcal{C} = \operatorname{dom}(h|_{\mathcal{C}}) = \{x \in \mathcal{C} : h|_{\mathcal{C}}(x) < \infty\}$. Denote by $T^* \subseteq \mathcal{V}^*$ the dual space of T. Following Rockafellar [39, Chap. 26], recall that the *subdifferential* at $x \in \mathcal{C}$ of $h|_{\mathcal{C}}$ is

$$\partial(h|_{\mathcal{C}})(x) = \{v \in T^* : h(x') \geq h(x) + \langle v, x' - x\rangle \ \text{ for all } x' \in \mathcal{C}\}, \quad \text{for all } x \in \mathcal{C}\,. \tag{B.3}$$

Any $v \in \partial(h|_{\mathcal{C}})(x)$ is called a *subgradient* of $h|_{\mathcal{C}}$ at x; note that subgradients belong to T^*, that is, they are defined only by how they act on elements of T. The *prox-domain $\mathcal{C}_h$ of h on $\mathcal{C}$* is the domain of subdifferentiability of $h|_{\mathcal{C}}$, that is,

$$\mathcal{C}_h = \operatorname{dom}\partial(h|_{\mathcal{C}}) = \{x \in \mathcal{C} : \partial(h|_{\mathcal{C}})(x) \neq \varnothing\}\,. \tag{B.4}$$

By standard results [39, Theorem 26.1], we have $\operatorname{ri}\mathcal{C} \subseteq \mathcal{C}_h \subseteq \mathcal{C}$.

Next, we introduce a way to continuously select one of the subgradients of $h|_{\mathcal{C}}$ at $x \in \mathcal{C}_h$: A *continuous selection of $\partial(h|_{\mathcal{C}})$ on $\mathcal{C}_h$* is a *continuous* map $\nabla(h|_{\mathcal{C}})\colon \mathcal{C}_h \to T^* \subseteq \mathcal{V}^*$ such that $\nabla(h|_{\mathcal{C}})(x) \in \partial(h|_{\mathcal{C}})(x)$ for all $x \in \mathcal{C}_h$.[4]

We suppress here on the subscript $|_{\mathcal{C}}$, with the understanding that $\partial h(x)$ is defined as in Eq. (B.3) (and does not denote the subdifferential of the ambient function h, that would be a singleton whenever h is differentiable); and that $\nabla h\colon \mathcal{C}_h \to T^* \subseteq \mathcal{V}^*$ is a continuous selection of ∂h on $\mathcal{C}_h$.

Definition 2. Let $\mathcal{V}$ be a normed vector space with norm $\|\cdot\|$, let $\mathcal{V}^*$ denote its dual, and let $\mathcal{C} \subseteq \mathcal{V}$ be closed and convex. A convex function $h\colon \mathcal{V} \to \mathbb{R} \cup \{\infty\}$ is called a *Bregman regularizer* on $\mathcal{C}$ if:

(*a*) h is finite and continuous on $\mathcal{C}$.

(*b*) The subdifferential of h admits a continuous selection ∇h on the *prox-domain* $\mathcal{C}_h \coloneqq \operatorname{dom}\partial h$.

[4] If $h|_{\mathcal{C}}$ is differentiable on $\operatorname{ri}\mathcal{C}$—which is the case for all of the commonly employed regularizers—a standard choice is to take the continuous selection $\nabla(h|_{\mathcal{C}})$ equal to the gradient of $h|_{\mathcal{C}}$ on $\operatorname{ri}\mathcal{C}$—that is, the pull-back of ∇h to $\operatorname{ri}\mathcal{C}$—and to extend it analytically to $\mathcal{C}_h \setminus \operatorname{ri}\mathcal{C}$.

(*c*) h is *strongly convex* on $\mathcal{C}$, that is, there exists some $\kappa > 0$ such that

$$h(x') \geq h(x) + \langle \nabla h(x), x' - x \rangle + \frac{\kappa}{2}\|x' - x\|^2 \qquad \text{for all } x \in \operatorname{dom} \partial h,\, x' \in \mathcal{C}\,. \tag{B.5}$$

We will say that h is *steep* if $\mathcal{C}_h = \operatorname{ri} \mathcal{C}$. We also write[5]

$$h^*(y) = \max_{x \in \mathcal{C}} \{\langle y, x \rangle - h(x)\} \qquad \text{for all } y \in \mathcal{V}^*. \tag{B.6}$$

for the convex conjugate of h, and define:

(*a*) The *Bregman divergence* of h:

$$D(p, x) = h(p) - h(x) - \langle \nabla h(x), p - x \rangle \qquad \text{for all } p \in \mathcal{C},\, x \in \mathcal{C}_h. \tag{B.7a}$$

(*b*) The associated *prox-mapping:*

$$P_x(v) = \arg\min_{u \in \mathcal{C}} \{\langle v, x - u \rangle + D(u, x)\} \qquad \text{for all } x \in \mathcal{C}_h,\, v \in \mathcal{V}^*. \tag{B.7b}$$

(*c*) The induced *Fenchel coupling:*

$$F(p, y) = h(p) + h^*(y) - \langle y, p \rangle \qquad \text{for all } p \in \mathcal{X}_h,\, y \in \mathcal{Y}. \tag{B.7c}$$

(*d*) The (primal-dual) *mirror*—or *regularized choice—map* of h:

$$Q(y) = \arg\max_{x \in \mathcal{C}} \{\langle y, x \rangle - h(x)\} \qquad \text{for all } y \in \mathcal{Y}. \tag{B.7d}$$

In the rest of this appendix, we collect several properties of the objects introduced in Definition 2.

Proposition B.1 (Mirror and proximal map)**.** *Let h be a κ-strongly convex regularizer on $\mathcal{C}$. Then:*

1. *The mirror map $Q\colon \mathcal{V}^* \to \mathcal{C}$ of Eq. (B.7d) is single valued on $\mathcal{C}$, and generically not injective. In particular, for all $x \in \operatorname{dom} \partial h$ and $y, y' \in \mathcal{V}^*$, we have that*

$$Q(y) = Q(y') = x \iff y \in \partial h(x) \text{ and } y' \in \partial h(x)\,, \tag{B.8a}$$

$$Q(y) = Q(y + v) \quad \text{for all } y \in \mathcal{V}^* \text{ and } v \in \mathrm{NC}_{Q(y)}\,. \tag{B.8b}$$

2. *Q is generically not surjective: precisely, we have $\operatorname{ri} \mathcal{C} \subseteq \operatorname{im} Q = \mathcal{C}_h \subseteq \mathcal{C}$. Furthermore, if h is steep, then $\operatorname{ri} \mathcal{C} \supseteq \operatorname{im} Q$.*

3. *The convex conjugate $h^*\colon \mathcal{V}^* \to \mathbb{R}$ is differentiable, and*

$$Q(y) = \nabla h^*(y) \quad \text{for all } y \in \mathcal{V}^*. \tag{B.9}$$

4. *Q is $1/\kappa$-Lipschitz continuous.*

5. *If h admits a continuous selection $\nabla h\colon \mathcal{C}_h \to T^* \subseteq \mathcal{V}^*$ of ∂h on $\mathcal{C}$, then*

$$\partial h(x) = \nabla h(x) + \mathrm{NC}_x \quad \text{for all } x \in \mathcal{C}_h\,. \tag{B.10}$$

6. *The* proximal map (B.7b) *$P\colon \mathcal{C}_h \times \mathcal{V}^* \to \mathcal{C}$ induced by h is equivalent to*

$$P_x(v) = Q(\nabla h(x) + v) \qquad \text{for all } x \in \mathcal{C}_h,\, v \in \mathcal{V}^*\,. \tag{B.11}$$

Furthermore, if $x \in \operatorname{ri} \mathcal{C}$, then

$$P_x(v) = Q(y + v) \quad \text{for any } y \in \partial h(x) \text{ and all } v \in \mathcal{V}^*\,. \tag{B.12}$$

7. *Fix some $y, y' \in \mathcal{V}^*$, and let $x' = Q(y')$. Then, for all $p \in \mathcal{C}$,*

$$\langle y' - y, p - x' \rangle \leq \langle \nabla h(x') - y, p - x' \rangle\,. \tag{B.13}$$

We give a technical lemma, and proceed to the proof of Proposition B.1.

[5]The terminology "steep" alludes to the fact that $\nabla h(x)$ grows unbounded near the boundary $\operatorname{bd} \mathcal{C}$ of $\mathcal{C}$.

Lemma B.2. *Fix $x \in \mathcal{C}_h$, and let $y \in \partial h(x)$. Then*

a. $y + v \in \partial h(x)$ *for all* $v \in \mathrm{NC}_x$. (B.14a)

b. $y \pm n \in \partial h(x)$ *for all* $n \in T^\circ$. (B.14b)

Proof. For the first point, by definitions, $h(x) + \langle y, p - x\rangle + \langle v, p - x\rangle \le h(p)$ for all $p \in \mathcal{C}$; and similarly for the second. ∎

Remark. The $\pm$ in Eq. (B.14b) is redundant, since T° is a linear space; we maintain it as a visual reminder that, if $x \in \operatorname{dom} \partial h$, $y \in \partial h(x)$, and $v \in \mathrm{NC}_x$, then in general $y - v \notin \partial h(x)$.

Proof of Proposition B.1. Since these results are relatively standard, we provide only proof sketches or precise references to the literature.

1. By strong convexity of h, the maximum in Eq. (B.7d) is attained for all $y \in \mathcal{V}^*$; and x is optimal if and only if $y \in \partial h(x)$. Eq. (B.8b) follows immediately by Eq. (B.8a) and Lemma B.2.

2. By Eq. (B.8a) it follows that $\operatorname{im} Q = \mathcal{X}_h$, and the "steep" part of the assertion follows Definition 2. The rest of the statements follows by standard arguments, cf. e.g., [39, Ch. 26].

3. See [10, Proposition 5.4.8, Appendix B].

4. See [41, Theorem 12.60(b)].

5. For the "$\subseteq$" part, fix $x \in \mathcal{C}_h$, and let $y \in \partial h(x)$; we want to show that $y - \nabla(x) \in \mathrm{NC}_x$. We consider the case in which $x \in \operatorname{ri} \mathcal{C}$; for the case $x \in \mathcal{C}_h \setminus \operatorname{ri} \mathcal{C}$, see e.g., [24, Prop. B.1 (e)]. Take any $p \in \mathcal{C}$, and let $d = p - x$. For ϵ small enough, by definition of continuous selection and of subgradient, a first-order expansion yields $h(x+\epsilon d) = h(x)+\epsilon\langle \nabla h(x), d\rangle + \mathrm{o}(\epsilon) \ge h(x)+\epsilon\langle y, d\rangle$. Dividing both sides by ϵ and letting $\epsilon \to 0$ then gives $\langle y - \nabla h(x), d\rangle \le 0$.

   For the "$\supseteq$" part, fix $x \in \mathcal{C}_h$; we want to show that $\nabla h(x) + v \in \partial h(x)$ for all $v \in \mathrm{NC}_x$, which follows by Eq. (B.14a) since $\nabla h(x) \in \partial h(x)$.

6. For the first part of the statement,

$$\begin{aligned} P_x(v) &= \operatorname*{arg\,min}_{u \in \mathcal{C}} \{\langle v, x - u\rangle + D(u, x)\} \\ &= \operatorname*{arg\,max}_{u \in \mathcal{C}} \{\langle \nabla h(x) + v, u\rangle - h(u)\} \\ &= Q(\nabla h(x) + v)\,. \end{aligned} \tag{B.15}$$

   For the second: fix $x \in \mathcal{C}_h$; any $y \in \partial h(x)$ can be written as $y = \nabla h(x) + n$ for some $n \in \mathrm{NC}_x$. Hence, for any $v \in \mathcal{V}^*$, we have $P_x(v) = Q(\nabla h(x) + v) = Q(y + v - n)$, and we can not say anything more. However, if $x \in \operatorname{ri} \mathcal{C}$, then $n \in T^\circ$ by Lemma B.1, so it follows by Eq. (B.14b) that $y + v - n$ and $y + v$ are subgradients at the same point, that is, $Q(y + v - n) = Q(y + v)$. [Equivalently: if $x \in \operatorname{ri} \mathcal{C}$ then $n \in T^\circ$, hence $-n \in T^\circ$ as well, and $Q(y + v - n) = Q(y + v)$ by B.8b].

7. By Eq. (B.10), $y' = \nabla h(x') + v$ for some $v \in \mathrm{NC}_x$, hence $\langle y' - \nabla h(x'), p - x'\rangle \le 0$ for all $p \in \mathcal{C}$; the conclusion follows subtracting $\langle y, p - x'\rangle$ from both sides. ∎

Proposition B.2 (Fenchel coupling)**.** *Let h be a κ-strongly convex regularizer on $\mathcal{C}$. Then, for all $p \in \mathcal{C}$ and all $y, y' \in \mathcal{V}^*$, we have*[6]

1. $F(p, y') = F(p, y) + F(Q(y), y') + \langle y' - y, Q(y) - p\rangle\,.$ (B.16a)

2. $F(p, y) \ge 0$ *with equality if and only if* $p = Q(y)$. (B.16b)

3. $F(p, y) \ge \frac{\kappa}{2} \|Q(y) - p\|^2$. (B.16c)

Proof.

[6] Eq. (B.16a) is a primal-dual version of the standard "three-point identity" of Chen & Teboulle [13].

1. Eq. (B.16a) follows by summing the Fenchel coupling terms, replacing their definition, and noting that $h^*(y) = \langle y, Q(y)\rangle - h(Q(y))$.

2. By the Fenchel-Young inequality, $h(p) + h^*(y) - \langle p, y\rangle \geq 0$ for all $p \in \mathcal{C}$ and all $y \in \mathcal{V}^*$, with equality if and only if $y \in \partial h(p)$, which is equivalent to $p = Q(y)$ by Eq. (B.8a).

3. Follows by strong convexity of h; see e.g., [24, Prop. B.2].

■

Lemma B.3 (Fenchel-Young inequality in Peter-Paul form)**.** *For any $x \in \mathcal{V}$, $y \in \mathcal{V}^*$, and $k > 0$,*

$$\langle x, y\rangle \leq \frac{k}{2}\|x\|^2 + \frac{1}{2k}\|y\|_*^2 . \tag{B.17}$$

Proof. Immediate from the Fenchel-Young inequality applied to the function $x \mapsto (k/2)\|x\|^2$. ■

C Algorithm analysis

For notational clarity, we suppress from the proofs in this appendix the player index i, so that the algorithmic template of (FTRL+) reads

$$\begin{aligned} &a)\ \textit{Extrapolation:} && y_{t+1/2} = y_t + \gamma_t \hat{v}_t && x_{t+1/2} = Q(\eta_{t+1/2}\, y_{t+1/2}) , \\ &b)\ \textit{Update:} && y_{t+1} = y_t + \gamma_{t+1/2} v_{t+1/2} && x_{t+1} = Q(\eta_{t+1}\, y_{t+1}) . \end{aligned} \tag{C.1}$$

Note that the primal steps of (FTRL+) can be written as

$$\begin{aligned} &a)\ \textit{Extrapolation:} && x_{t+1/2} = Q(\eta_{t+1/2} y_t + \eta_{t+1/2}\gamma_t \hat{v}_t) , \\ &b)\ \textit{Update:} && x_{t+1} = Q(\eta_{t+1} y_t + \eta_{t+1}\gamma_{t+1/2} v_{t+1/2}) . \end{aligned} \tag{C.2}$$

We recall that the player's strategy space $\mathcal{X}$ is a convex subset of a real vector space $\mathcal{V}$. As per Definition 2 in Appendix B, we denote by $\mathcal{Y} = \mathcal{V}^*$ the dual space of $\mathcal{V}$, and by $Q\colon \mathcal{Y} \to \mathcal{X}$ (resp. $P\colon \mathcal{X}_h \times \mathcal{Y} \to \mathcal{X}$) the mirror (resp. proximal) map induced by a κ-strongly convex regularizer $h\colon \mathcal{X} \to \mathbb{R}$. We also write $v_t \equiv v(x_t)$ and $v_{t+1/2} \equiv v(x_{t+1/2})$; with this convention, the payoff signals defined in Eq. (6) read $\hat{v}_{t+1/2} = v_{t+1/2}$ and $\hat{v}_t = \lambda v_t + (1-\lambda) v_{t-1/2}$, where $\lambda \in [0, 1]$.

C.1. Mirror descent and its friends. In a nutshell, mirror descent replaces gradient steps in standard gradient methods by a "mirror step"—a non-Euclidean projection step which is tailored to the geometry of the problem under study. Concretely, the *mirror descent* (MD) algorithm unfolds as

$$x_{t+1} = P_{x_t}(\gamma_t v_t) , \tag{MD}$$

for $t = 1, 2, \dots$, with P_{x_t} the *prox-mapping* operator (see Definition 2 in Appendix B); $\gamma_t > 0$ a variable *step-size* sequence; and $v_t \in \mathcal{V}^*$ a sequence of dual "payoff-like" vectors. If γ_t is decreasing, (MD) exhibits the somewhat strange feature that "new information enters the algorithm with decreasing weights", i.e., more recent information seem to be discounted relative to older data. Nesterov [33] leveraged this observation as the starting point for the *dual averaging* (DA) scheme

$$y_{t+1} = y_t + \gamma_t v_t \qquad x_{t+1} = Q(\eta_{t+1} y_{t+1}) , \tag{DA$_{\gamma,\eta}$}$$

where Q is the *mirror map* (see Definition 2 in Appendix B). (DA$_{\gamma,\eta}$) includes both a step-size sequence $\gamma_t > 0$ *pre*-multiplying new gradients, and a *learning rate* sequence $\eta_t > 0$ *post*-multiplying the aggregation of information $y_{t+1} \leftarrow y_t + \gamma_t v_t$ in the dual—hence the name, see also [47].

Instances of (DA$_{\gamma,\eta}$) with constant learning rate $\eta_t \leftarrow 1$ [which goes by the name of *lazy mirror descent* (LMD)] or with constant step-size $\gamma_t \leftarrow 1$ are closely related to the *Follow-the-Regularized-Leader* (FTRL) online learning template [42, 43], viz.

$$x_{t+1} = \arg\max_{x\in\mathcal{X}} \left\{ \sum_{s=1}^{t} \langle v_s, x\rangle - \lambda\, h(x) \right\} \tag{FTRL}$$

for some *regularization parameter* $\lambda > 0$ and *Bregman regularizer* h (see Definition 2 in Appendix B). Clearly, these three methods—(LMD), (DA) with constant step size, and (FTRL)—are equivalent when $\gamma_t = \eta_t = 1/\lambda$; however, beyond the constant parameter case, there is some ambiguity in the literature concerning these methods. [Adding to the confusion, (MD) and (LMD) are also equivalent when h is steep; for completeness, we prove this fact below.] To connect with this trend in the literature, we will also use the term "FTRL" for the general template (DA$_{\gamma,\eta}$); having said that, we stress that different parameter schedules can lead to drastically different behaviors.

Equivalence of MD and LMD To begin, we revisit the well-known equivalence of *mirror descent* (MD) and *lazy mirror descent* (LMD) when h is steep; the same reasoning will then be employed to establish the relation between (FTRL+) and several existing schemes under steep regularization. Recall from Section 4.1 that the basic update step of (LMD) can be written as

$$y^+ = y + v \quad \text{and} \quad x^+ = Q(y^+) \tag{C.3}$$

while the corresponding update step for (MD) is

$$x^+ = P_x(v) \tag{C.4}$$

for some $y, v \in \mathcal{V}^*$. With this in mind, the main ingredient of this equivalence is Eq. (B.12) — which, if h is steep, holds for all $y \in \mathcal{V}^*$, $x = Q(y)$, and $v \in \mathcal{V}^*$.

Corollary 1. *Suppose that* (MD) *and* (LMD) *are initialized respectively at $x_1 \in \mathcal{X}_h$ and $y_1 \in \mathcal{V}^*$, with $x_1 = Q(y_1)$. If h is steep and the algorithms are run with the same input sequence v_t, $t = 1, 2, \dots$, then they generate the same sequence of iterates.*

Proof. If h is steep then Eq. (B.12) holds true; this implies the equivalence of Eqs. (C.3) and (C.4). ∎

C.2. Relation between (FTRL+) and other extrapolation methods. For reference, we provide in this appendix some representative examples (all stated with h assumed steep and player indices suppressed for simplicity); Lemma C.1 details how each of these algorithms arises as a special case of (FTRL+) under suitable parameter choices.

1. *Mirror-prox* (MP) [6, 31] updates as
$$x_{t+1/2} = P_{x_t}(\alpha_t\, v(x_t)) \qquad x_{t+1} = P_{x_t}(\alpha_t\, v(x_{t+1/2}))\,. \tag{MP}$$
2. *Dual extrapolation* (DEX) [32] updates as
$$x_{t+1/2} = P_{x_t}(\beta_t\, v(x_t)) \qquad y_{t+1} = y_t + v(x_{t+1/2}) \qquad x_{t+1} = Q(\beta_{t+1}\, y_{t+1})\,. \tag{DEX}$$
3. *Optimistic FTRL* (OptFTRL) [2, 30, 44] updates as
$$x_{t+1/2} = Q\left(\frac{1}{\lambda_t}\left(\textstyle\sum_{s=1}^{t} v(x_{s-1/2}) + v(x_{t-1/2})\right)\right). \tag{OptFTRL}$$
4. *Optimistic dual averaging* (OptDA) [22] updates as
$$x_{t+1/2} = P_{x_t}(\beta_t\, v(x_{t-1/2})) \qquad y_{t+1} = y_t + v(x_{t+1/2}) \qquad x_{t+1} = Q(\beta_{t+1}\, y_{t+1})\,. \tag{OptDA}$$

Lemma C.1. *With a steep regularizer, the primal sequences generated by* (FTRL+) *are equivalent to those generated by*

- (DEX) *for $\gamma_t = \gamma_{t+1/2} = 1$, $\eta_t = \eta_{t+1/2} = \beta_t$, and $\hat{v}_t = v_t$;*
- (OptDA) *for $\gamma_t = \gamma_{t+1/2} = 1$, $\eta_t = \eta_{t+1/2} = \beta_t$, and $\hat{v}_t = v_{t-1/2}$;*
- (MP) *for $\gamma_t = \gamma_{t+1/2} = \alpha_t$, $\eta_t = \eta_{t+1/2} = 1$, and $\hat{v}_t = v_t$;*
- (OptFTRL) *for $\gamma_t = \gamma_{t+1/2} = 1$, arbitrary $\eta_t > 0$, $\eta_{t+1/2} = 1/\lambda_t$, $\hat{v}_t = v_{t-1/2}$, and under* (Init).

Proof of Lemma C.1. If the regularizer h is steep, then the iterates x_t and $x_{t+1/2}$ produced by all the methods belong to the relative interior of $\mathcal{X}$ for all $t = 1, 2, \dots$ by Item 2 of Proposition B.1 and Eq. (B.11), which allows us to repeatedly invoke Eq. (B.12).

- (DEX), (OptDA) : We treat these cases together, since they differ only in the form of the signal $\hat{v}_t$. We need to show that, if $\gamma_{t+1/2} = \gamma_t = 1$, $\eta_{t+1/2} = \eta_t = \beta_t$, then the update rule (C.1) is equivalent to
$$\begin{aligned} &a)\ \textit{Extrapolation:} && x_{t+1/2} = P_{x_t}(\beta_t\, \hat{v}_t) \\ &b)\ \textit{Update:} \quad y_{t+1} = y_t + v_{t+1/2} && x_{t+1} = Q(\beta_{t+1}\, y_{t+1}) \end{aligned} \tag{DEX-OptDA}$$
By Eq. (B.8a), $\beta_t y_t \in \partial h(x_t)$, hence invoking Eq. (B.12) we see that the primal extrapolation step of (DEX-OptDA) is $x_{t+1/2} = Q(\beta_t y_t + \beta_t \hat{v}_t)$. Thus, $\eta_{t+1/2} = \beta_t$ and $\gamma_t = 1$ ensure that the extrapolation phases of the two algorithms are equivalent — cf. Eq. (C.2). The update phases are clearly equivalent under the choice $\gamma_{t+1/2} = 1$ and $\eta_t = \beta_t$.

- (MP): Analogously to the previous case, the primal steps of (MP) read $x_{t+1/2} = Q(y_t + \alpha_t v_t)$ and $x_{t+1} = Q(y_t + \alpha_t v_{t+1/2})$, for any $y_t \in \partial h(x_t)$, i.e., for any y_t such that $x_t = Q(y_t)$.
- (OptFTRL): Instantiating Eq. (C.1) with the given parameters combination yields $x_{t+1/2} = Q\big(1/\lambda_t \big(y_t + v_{t-1/2}\big)\big)$, and telescoping the dual update step of Eq. (C.1) with initialization condition (Init) gives $y_t = \sum_{s=1}^{t} v_{s-1/2}$. ∎

C.3. Template inequality for (FTRL+). In this section we prove the template inequality of Proposition 1, which we restate below for convenience. For the analysis to come it is convenient to define the *difference of gradients*

$$g_{i,t} \coloneqq v_{i,t+1/2} - \hat{v}_{i,t}\,, \tag{C.5}$$

so that the adaptive policy (12) reads

$$\eta_{i,t} = \frac{1}{\sqrt{1 + \sum_{s=1}^{t-1} \|g_{i,s}\|_*^2}}\,. \tag{C.6}$$

Proposition 1 (Template inequality)**.** *Suppose that each player is running* (FTRL+) *with payoff signals as per* (6)*. Then, for all $p_i \in \mathcal{X}_i$, $i \in \mathcal{N}$, and all $t = 1, 2, \dots$, we have:*

$$E_{i,t+1} \leq E_{i,t} + \gamma_{i,t+1/2} \langle \hat{v}_{i,t+1/2}, x_{i,t+1/2} - p_i \rangle \tag{9a}$$

$$+ \left(\frac{1}{\eta_{i,t+1}} - \frac{1}{\eta_{i,t}} \right) \cdot [h_i(p_i) - \min h_i] \tag{9b}$$

$$+ \gamma_{i,t+1/2} \langle \hat{v}_{i,t+1/2} - \hat{v}_{i,t}, x_{i,t+1} - x_{i,t+1/2} \rangle \tag{9c}$$

$$- \frac{\gamma_{i,t+1/2}}{\gamma_{i,t}} \left[\frac{F_i(x_{i,t+1}, \eta_{i,t+1/2} y_{i,t+1/2})}{\eta_{i,t+1/2}} + \frac{F_i(x_{i,t+1/2}, \eta_{i,t+1/2} y_{i,t})}{\eta_{i,t+1/2}} \right] \tag{9d}$$

$$+ \frac{\gamma_{i,t+1/2}}{\gamma_{i,t}} \frac{F_i(x_{i,t+1}, \eta_{i,t+1/2} y_{i,t})}{\eta_{i,t+1/2}} - \frac{F_i(x_{i,t+1}, \eta_{i,t} y_{i,t})}{\eta_{i,t}}\,. \tag{9e}$$

Proof. We want to show that

$$\begin{aligned}
E_{i,t+1} \leq\ & E_{i,t} && \text{(energy step)}\\
& - \gamma_{i,t+1/2} \langle v_{i,t+1/2}, p_i - x_{i,t+1/2} \rangle && \text{(regret)}\\
& + \left(\frac{1}{\eta_{i,t+1}} - \frac{1}{\eta_{i,t}} \right) \cdot [h_i(p_i) - \min h_i] && \text{(regularizer depth)}\\
& + \gamma_{i,t+1/2} \langle g_{i,t}, x_{i,t+1} - x_{i,t+1/2} \rangle && \text{(difference of gradients)}\\
& - \frac{\gamma_{i,t+1/2}}{\gamma_{i,t}} \left[\frac{F_i(x_{i,t+1}, \eta_{i,t+1/2} y_{i,t+1/2})}{\eta_{i,t+1/2}} + \frac{F_i(x_{i,t+1/2}, \eta_{i,t+1/2} y_{i,t})}{\eta_{i,t+1/2}} \right] && \text{(triangle)}\\
& + \frac{\gamma_{i,t+1/2}}{\gamma_{i,t}} \frac{F_i(x_{i,t+1}, \eta_{i,t+1/2} y_{i,t})}{\eta_{i,t+1/2}} - \frac{F_i(x_{i,t+1}, \eta_{i,t} y_{i,t})}{\eta_{i,t}}\,, && \text{(coupling)}
\end{aligned}$$

with Eq. (C.5) defining the difference of gradients $g_{i,t}$.

To begin observe that, for any $v, \hat{v} \in \mathcal{Y}$ and $p, x, x' \in \mathcal{V}$, the following identity holds:

$$\langle v, p - x \rangle \equiv \langle v, p - x' \rangle + \langle \hat{v}, x' - x \rangle + \langle v - \hat{v}, x' - x \rangle\,. \tag{C.7}$$

Consider the update given by Eq. (C.1). For any $p \in \mathcal{X}$ and $t = 1, 2, \dots$, let $g_t = v_{t+1/2} - \hat{v}_t$ and instantiate the identity above with $v \leftarrow v_{t+1/2}$, $\hat{v} \leftarrow \hat{v}_t$, $x \leftarrow x_{t+1/2}$, $x' \leftarrow x_{t+1}$ to get

$$\langle v_{t+1/2}, p - x_{t+1/2} \rangle \equiv \underbrace{\langle v_{t+1/2}, p - x_{t+1} \rangle}_{\text{update}} + \underbrace{\langle \hat{v}_t, x_{t+1} - x_{t+1/2} \rangle}_{\text{extrapolation}} + \underbrace{\langle g_t, x_{t+1} - x_{t+1/2} \rangle}_{\text{difference of gradients}}\,. \tag{C.8}$$

Next, bound the "update" (resp. "extrapolation") term of Eq. (C.8) using the "update" (resp. "extrapolation") phase of Eq. (C.1).

Step 1: Bound update term. The dual step in the update phase of Eq. (C.1) reads $\gamma_{t+1/2} v_{t+1/2} = y_{t+1} - y_t$, thus we get

$$\begin{aligned}
\gamma_{t+1/2}\langle v_{t+1/2}, p - x_{t+1}\rangle &\equiv \left\langle \frac{\eta_{t+1}}{\eta_{t+1}} y_{t+1} - \frac{\eta_t}{\eta_t} y_t, p - x_{t+1} \right\rangle \pm \frac{\eta_{t+1}}{\eta_t}\langle y_{t+1}, p - x_{t+1}\rangle \\
&= \left[\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t}\right]\langle \eta_{t+1} y_{t+1} - 0, p - x_{t+1}\rangle + \frac{1}{\eta_t}\langle p - x_{t+1}, \eta_{t+1} y_{t+1} - \eta_t y_t\rangle \\
&= \left[\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t}\right][F(p,0) - F(p, \eta_{t+1} y_{t+1}) - F(x_{t+1}, 0)] \\
&\quad + \frac{1}{\eta_t}[F(p, \eta_t y_t) - F(p, \eta_{t+1} y_{t+1}) - F(x_{t+1}, \eta_t y_t)] \qquad \% \text{ Eq. (B.16a)} \\
&= E_t - E_{t+1} + \left[\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t}\right](h(p) - h(x_{t+1})) - \frac{F(x_{t+1}, \eta_t y_t)}{\eta_t} \\
&\leq E_t - E_{t+1} + \left[\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t}\right](h(p) - \min h) - \frac{F(x_{t+1}, \eta_t y_t)}{\eta_t}, \qquad \text{(C.9)}
\end{aligned}$$

where $E_t = {}^{F(p, \eta_t y_t)}/_{\eta_t}$ as per Eq. (8); in the last step we used the fact that η_t is non-increasing.

Step 2: Bound extrapolation term. The dual step in the extrapolation phase of Eq. (C.1) is $y_{t+1/2} = y_t + \gamma_t \frac{\eta_{t+1/2}}{\eta_{t+1/2}} \hat{v}_t$, thus by the three-point identity (B.16a) we get

$$\begin{aligned}
\gamma_t \eta_{t+1/2}\langle \hat{v}_t, x_{t+1} - x_{t+1/2}\rangle &= \langle \eta_{t+1/2} y_{t+1/2} - \eta_{t+1/2} y_t, x_{t+1} - x_{t+1/2}\rangle \\
&= F(x_{t+1}, \eta_{t+1/2} y_t) - F(x_{t+1}, \eta_{t+1/2} y_{t+1/2}) - F(x_{t+1/2}, \eta_{t+1/2} y_t)\,. \qquad \text{(C.10)}
\end{aligned}$$

To conclude, replace the bounds obtained in Steps 1 and 2 into Eq. (C.8):

$$\begin{aligned}
\langle v_{t+1/2}, p - x_{t+1/2}\rangle \leq \frac{1}{\gamma_{t+1/2}}&\left[E_t - E_{t+1} + \left[\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t}\right](h(p) - \min h) - \frac{F(x_{t+1}, \eta_t y_t)}{\eta_t}\right] \\
&+ \frac{1}{\gamma_t \eta_{t+1/2}}\left[F(x_{t+1}, \eta_{t+1/2} y_t) - F(x_{t+1}, \eta_{t+1/2} y_{t+1/2}) - F(x_{t+1/2}, \eta_{t+1/2} y_t)\right] \\
&+ \langle g_t, x_{t+1} - x_{t+1/2}\rangle\,. \qquad \text{(C.11)}
\end{aligned}$$

A minor reorganization, together with the reinstatement of the player index, yields the inequality as stated in Proposition 1. ∎

D Regret analysis

D.1. Adversarial regret analysis. In this appendix we prove the first part of Theorem 4, which we restate below for convenience; the second part is proved in Appendix D.2.

Theorem 4 (Individual regret guarantees)**.** *Suppose that player $i \in \mathcal{N}$ is running* (FTRL+) *with payoff signals as per* (6)*, parameters as per* (10) *and* (12)*, and $\gamma_i = 1$. Then, letting $H_i := \max h_i - \min h_i$ and $G_i = \sup_{x \in \mathcal{X}} \|v_i(x)\|_*$, we have:*

(a) The regret of player $i \in \mathcal{N}$ against any sequence $X_{-i,t} \in \mathcal{X}_{-i}$, $t = 1, 2, \dots$, is bounded as

$$\mathrm{Reg}_i(T) \leq (H_i + 1/\kappa_i)(1 + 4G_i^2 T)^{1/2} + 2\,G_i^2/\kappa_i = \mathcal{O}\big(\sqrt{T}\big)\,. \qquad \text{(14a)}$$

(b) If, in addition, all players follow (FTRL+) *as above and the game is harmonic, then*

$$\mathrm{Reg}_i(T) = \mathcal{O}(H_i + 1/\kappa_i)\,. \qquad \text{(14b)}$$

Remark 1. For any $p \in \mathcal{X}$, we write for brevity $\underline{h}(p) := h(p) - \min h$. It follows immediately by Eq. (B.7c) that $\underline{h}(p) = F(p, 0)$. As in the previous section, we carry out the proof suppressing the player index.

Remark 2. In the following we will refer to the terms of the template inequality of Proposition 1 by the labels introduced in the corresponding proof in Appendix C.3.

Proof of Theorem 4–(a). We begin by considering the general update scheme of (FTRL+)—which includes all four parameters $\gamma_t, \eta_t, \gamma_{t+1/2}, \eta_{t+1/2}$. Observe that, for each player $i \in \mathcal{N}$, the coupling term in the template inequality of Proposition 1 is non-positive if

$$\eta_{i,t} = \eta_{i,t+1/2} \quad \text{and} \quad \gamma_{i,t+1/2} \leq \gamma_{i,t} \,. \tag{D.1}$$

In particular, coupling vanishes if and only if Eq. (D.1) holds with equality. We begin our analysis with the parameters constrained by Eq. (D.1), and specialize along the way to the setting of Theorem 4, with parameters constrained by Eq. (10).

For any $p \in \mathcal{X}$, consider the template inequality of Proposition 1. By virtue of Eq. (D.1), the coupling term can be dropped from the bound, and we can write $\Delta_t \coloneqq \gamma_t - \gamma_{t+1/2} \geq 0$. Imposing the initialization condition Init, dropping also the second non-positive Fenchel addend of the triangle term, and summing over $t = 1, 2, \dots, T$, we obtain

$$\begin{aligned}
\mathrm{Reg}_p(T) &= \sum_{t=1}^{T} \langle v_{t+1/2}, p - x_{t+1/2} \rangle \\
&\leq \sum_{t=1}^{T} \frac{E_t}{\gamma_{t+1/2}} - \frac{E_{t+1}}{\gamma_{t+1/2}} + \sum_{t=1}^{T} \left[\frac{1}{\eta_{t+1}} - \frac{1}{\eta_t} \right] \frac{\underline{h}(p)}{\gamma_{t+1/2}} \\
&\quad + \sum_{t=1}^{T} \langle g_t, x_{t+1} - x_{t+1/2} \rangle - \frac{1}{\gamma_t} \frac{F(x_{t+1}, \eta_t y_{t+1/2})}{\eta_t} \\
&\leq \frac{E_1}{\gamma_{1+1/2}} - \frac{E_{T+1}}{\gamma_{T+1/2}} + \sum_{t=1}^{T-1} E_{t+1} \left(\frac{1}{\gamma_{t+1+1/2}} - \frac{1}{\gamma_{t+1/2}} \right) && \% \text{ telescope} \\
&\quad + \underline{h}(p) \left[\frac{1}{\eta_{T+1}\gamma_{T+1/2}} - \frac{1}{\eta_1 \gamma_{1+1/2}} \right] + \sum_{t=1}^{T-1} \frac{\underline{h}(p)}{\eta_{t+1}} \left(\frac{1}{\gamma_{t+1/2}} - \frac{1}{\gamma_{t+1+1/2}} \right) && \% \text{ telescope} \\
&\quad + \sum_{t=1}^{T} \frac{\kappa}{2\eta_t \gamma_t} \|x_{t+1} - x_{t+1/2}\|^2 + \frac{\eta_t \gamma_t}{2\kappa} \|g_t\|_*^2 - \frac{\kappa}{2\eta_t \gamma_t} \|x_{t+1} - x_{t+1/2}\|^2 \\
&&& \% \text{ Eqs. (B.17) and (B.16c)} \\
&\leq \frac{\underline{h}(p)}{\eta_{T+1}\gamma_{T+1/2}} + \sum_{t=1}^{T} \frac{\eta_t \gamma_{t+1/2}}{2\kappa} \|g_t\|_*^2 && \% \; E_1 - \underline{h}(p) = 0 \\
&\quad + \sum_{t=1}^{T-1} \left(E_{t+1} - \frac{\underline{h}(p)}{\eta_{t+1}} \right) \left(\frac{1}{\gamma_{t+1+1/2}} - \frac{1}{\gamma_{t+1/2}} \right) + \sum_{t=1}^{T} \frac{\eta_t \Delta_t}{2\kappa} \|g_t\|_*^2 \,. && \text{(D.2)}
\end{aligned}$$

Eq. (D.2) suggests choosing $\Delta_t = 0$, that is, setting $\gamma_{t+1/2} = \gamma_t$. Now if (FTRL+) were run with *decreasing* γ_t, the resulting bound would be subject to the negative result of Orabona & Pál [34, Th. 4], who showed that (MD) with a vanishing step-size can incur *superlinear* regret. Consequently, if (FTRL+) is to interpolate between convergence and regret guarantees, the most principled choice is to take a constant step-size $\gamma_t = \gamma > 0$—cf. the discussion following Eq. (10). This rules out the mirror-prox method (MP) of Nemirovski [31] and Auslender & Teboulle [6], a fact which we view as an important take-away of this reasoning.

Further imposing $\gamma = 1$ (as per the assumptions of Theorem 4), maximizing over $p \in \mathcal{X}$, and invoking the adaptive policy (12) for the learning rate η_t, we obtain the following bound for FTRL+:

$$\begin{aligned}
\mathrm{Reg}(T) &\leq \frac{\max h - \min h}{\eta_{T+1}} + \frac{1}{2\kappa} \sum_{t=1}^{T} \eta_t \|g_t\|_*^2 && \text{(D.3)} \\
&= \frac{H}{\eta_{T+1}} + \frac{1}{2\kappa} \sum_{t=1}^{T} \eta_{t+1} \|g_t\|_*^2 + \frac{1}{2\kappa} \sum_{t=1}^{T} (\eta_t - \eta_{t+1}) \|g_t\|_*^2 \\
&\leq H \sqrt{1 + \sum_{t=1}^{T} \|g_t\|_*^2} + \frac{1}{\kappa} \sqrt{1 + \sum_{t=1}^{T} \|g_t\|_*^2} + \frac{1}{2\kappa} \sum_{t=1}^{T} (\eta_t - \eta_{t+1}) \|g_t\|_*^2 \,. \\
&&& \% \text{ Eq. (12) and Lemma D.2}
\end{aligned}$$

Recall from Eq. (C.5) that $g_t = v_{t+1/2} - \hat{v}_t$, and from Eq. (6b) that $\hat{v}_t = \lambda v_t + (1-\lambda) v_{t-1/2}$, so $\|g_t\|_* \leq \|v_{t+1/2}\|_* + \lambda \|v_t\|_* + (1-\lambda)\|v_{t-1/2}\|_* \leq 2G$. This leads to

$$\mathrm{Reg}(T) \leq \left(H + \frac{1}{\kappa}\right)\sqrt{1 + 4G^2 T} + \frac{2G^2}{\kappa}, \tag{D.4}$$

which concludes our proof. ∎

Remark 3. One could tune the numerator of Eq. (12) to the player-specific constant $\sqrt{H\kappa}$ and obtain a slight improvement over the bound of Theorem 4-(a). We set the numerator to 1 for ease of presentation, but the proof can be readily adapted to general values. One might also be tempted to set the addend +1 in the denominator of Eq. (12) to zero (and refrain from updating when $g_t = 0$ to avoid division by zero), thereby improving the bound and obtaining a scale-free adaptive policy. However, the addition of a positive constant in the denominator of Eq. (12)—while not essential for the proof of Theorem 4-(a)—is crucial for the proof of Theorem 4-(b). As before, this constant could be tuned to an optimal value with minor modifications to the argument; we set it equal to +1 for ease of presentation.

D.2. Analysis of regret in harmonic games. The crucial ingredient to prove Theorem 2, Theorem 3, and Theorem 4-(b) is the following

Proposition D.1 (Convergence of learning rate)**.** *Suppose that each player is running* (FTRL+) *in a harmonic game with payoff signals as per* (6) *and parameters as per* (10)*. Assume either of the following conditions holds:*

1. $\liminf_t \eta_{i,t} > 0$ *and* γ_i *fulfills Eq.* (11) $\big(\gamma_i \leq \min\{1, \kappa_i \cdot (\min_j m_j \min_j \kappa_j) \,/\, (16\, m_i N L_i^2)\}\big)$*, where* N *is the number of players in the game and, for each player,* κ_i *is the strong convexity modulus of the regularizer;* $m_i = \sum_{\alpha_i} \mu_{i\alpha_i}$ *is the mass of the harmonic measure* μ_i*; and* L_i *is the Lipschitz modulus of the player's payoff field.*

2. *Every player updates* $\eta_{i,t}$ *with the adaptive policy* (12) *and sets* $\gamma_{i,1} = 1$.

Then, for each $i \in \mathcal{N}$*, the sequence* $(\eta_{i,t})_{t\in\mathbb{N}}$ *converges to a positive constant* $\eta_{i,\infty}$*, and the sequence of difference of gradients squared* $\big(\|g_{i,t}\|_*^2\big)_{t\in\mathbb{N}}$ *defined by Eq.* (C.5) *is summable.*

We adapt the proof technique from [22]—who consider OptDA in the class of *variationally stable games*—to the setting of FTRL+ in harmonic games.

Remark 4. In the following we make repeated use of the fact that $\eta_{i,t}$ is bounded from above. Since it is monotonically decreasing, we set without loss of generality $\eta_{i,1} = 1$ for all $i \in \mathcal{N}$.

Remark 5. For all $i \in \mathcal{N}$ and any $p_i \in \mathcal{X}_i$, we write for brevity $\underline{h_i}(p_i) \coloneqq h_i(p_i) - \min h_i$. It follows immediately by Eq. (B.7c) that $\underline{h_i}(p_i) = F_i(p_i, 0)$.

Proof of Proposition D.1. The first part of the proof depends only on the template inequality of Proposition 1 under the choice of parameters of Eq. (10) and with signal Eq. (6b), and not on the specifics of the cases *(a)* and *(b)*; these will be treated separately in the second part of the proof.

For each $i \in \mathcal{N}$ and any $p_i \in \mathcal{X}_i$, sum Eq. (9) over $t = 1, 2, \dots, T$ to obtain the following bound for sequences generated by FTRL+:

$$\sum_{t=1}^{T} \langle v_{i,t+1/2}, p_i - x_{i,t+1/2}\rangle \leq \sum_{t=1}^{T} \frac{E_{i,t}}{\gamma_i} - \frac{E_{i,t+1}}{\gamma_i} + \sum_{t=1}^{T}\left[\frac{1}{\eta_{i,t+1}} - \frac{1}{\eta_{i,t}}\right]\frac{\underline{h_i}(p_i)}{\gamma_i} \tag{D.5a}$$

$$+ \sum_{t=1}^{T} \langle g_{i,t}, x_{i,t+1} - x_{i,t+1/2}\rangle \tag{D.5b}$$

$$- \sum_{t=1}^{T} \frac{F_i(x_{i,t+1}, \eta_{i,t} y_{i,t+1/2})}{\gamma_i \eta_{i,t}} - \sum_{t=1}^{T} \frac{F_i(x_{i,t+1/2}, \eta_{i,t} y_{i,t})}{\gamma_i \eta_{i,t}}. \tag{D.5c}$$

Observe by Eqs. (6b) and (C.5) that $g_{i,t} = v_{i,t+1/2} - \hat{v}_{i,t} = \big(v_{i,t+1/2} - v_{i,t}\big) + (1-\lambda_i)\big(v_{i,t} - v_{i,t-1/2}\big)$. As in the proof of the adversarial regret in Theorem 4, the index pattern of the difference between primal points in Eq. (D.5b) matches that of the first Fenchel term in Eq. (D.5c), so these can be

related by the Fenchel-Young inequality (B.17) and strong convexity (Eq. (B.16c)). Furthermore, after a minor reshuffling, the index pattern in the two Fenchel terms of Eq. (D.5c) can be arranged to match the index pattern of $g_{i,t}$: split the first Fenchel series of Eq. (D.5c) in two, and drop the last non-positive term from one half-series; split the second Fenchel series of Eq. (D.5c) in two, drop the first non-positive term from a half-series, and drop the whole second non-positive half-series; and telescope Eq. (D.5a), to get

$$\begin{aligned}\sum_{t=1}^{T}\langle v_{i,t+1/2}, p_i - x_{i,t+1/2}\rangle &\le \left(\frac{E_{i,1}}{\gamma_i} - \frac{\underline{h_i}(p_i)}{\gamma_i}\right) + \frac{H_i}{\gamma_i\eta_{i,T+1}} + \sum_{t=1}^{T}\langle g_{i,t}, x_{i,t+1} - x_{i,t+1/2}\rangle \\ &\quad - \left[\sum_{t=1}^{T}\frac{F_i(x_{i,t+1}, \eta_{i,t}y_{i,t+1/2})}{2\gamma_i\eta_{i,t}} + \sum_{t=1}^{T-1}\frac{F_i(x_{i,t+1}, \eta_{i,t}y_{i,t+1/2})}{2\gamma_i\eta_{i,t}}\right] \\ &\quad - \sum_{t=2}^{T}\frac{F_i(x_{i,t+1/2}, \eta_{i,t}y_{i,t})}{2\gamma_i\eta_{i,t}} \\ &\le \frac{H_i}{\gamma_i\eta_{i,T+1}} \qquad \% \;\; E_{i,1} - \underline{h_i}(p_i) = 0 \\ &\quad + \sum_{t=1}^{T}\left[\langle g_{i,t}, x_{i,t+1} - x_{i,t+1/2}\rangle - \frac{F_i(x_{i,t+1}, \eta_{i,t}y_{i,t+1/2})}{2\gamma_i\eta_{i,t}}\right] \\ &\quad - \sum_{t=2}^{T}\left[\frac{F_i(x_{i,t+1/2}, \eta_{i,t}y_{i,t})}{2} + \frac{F_i(x_{i,t}, \eta_{i,t-1}y_{i,t-1/2})}{2}\right]. \\ &\qquad \% \text{ Rename dummy index, } \eta_{i,t} \le 1, \gamma_i \le 1\end{aligned} \tag{D.6}$$

By the Fenchel-Young inequality and strong convexity (Eqs. (B.17) and (B.16c)), we conclude that

$$\begin{aligned}\sum_{t=1}^{T}\langle v_{i,t+1/2}, p_i - x_{i,t+1/2}\rangle &\le \frac{H_i}{\gamma_i\eta_{i,T+1}} + \frac{\gamma_i\|g_{i,1}\|_*^2}{\kappa_i} + \sum_{t=2}^{T}\frac{\gamma_i\eta_{i,t}}{\kappa_i}\|g_{i,t}\|_*^2 \\ &\quad - \frac{\kappa_i}{4}\sum_{t=2}^{T}\left(\|x_{i,t+1/2} - x_{i,t}\|^2 + \|x_{i,t} - x_{i,t-1/2}\|^2\right).\end{aligned} \tag{D.7}$$

The index pattern of the half-step sequences in the final line of Eq. (D.7) now matches that of $g_{i,t} = (v_{i,t+1/2} - v_{i,t}) + (1-\lambda_i)(v_{i,t} - v_{i,t-1/2})$. This correspondence, together with the Lipschitz continuity of the payoff fields, will allow us to bound the half-step sequences in (D.7) by invoking the following simple but crucial lemma:

Lemma D.1 (Difference of gradients)**.** *The difference of gradients squared defined by Eqs. (6b) and (C.5) which appears in the adaptive policy (12) is upper-bounded as*

$$\frac{\|g_{i,t}\|_*^2}{2} \le \|v_{i,t+1/2} - v_{i,t}\|_*^2 + \|v_{i,t} - v_{i,t-1/2}\|_*^2. \tag{D.8}$$

Proof. The claim follows (for any $\lambda_i \in [0,1]$) by Eqs. (6b) and (C.5), and the inequality $\|a+b\|^2 \le 2\|a\|^2 + 2\|b\|^2$. ∎

Remark. The extrapolative nature of FTRL+ is crucial for Lemma D.1 to hold. Setting $\hat{v}_{i,t} = 0$—which collapses FTRL+ to its vanilla variant—implies $g_{i,t} = v_{i,t} = v_{i,t+1/2}$ and $v_{i,t-1/2} = v_{i,t-1}$. This breaks the validity of Eq. (D.8), which in turn is essential for the rest of our argument.

To move on, fix a norm $\|\cdot\|_i$ for the ambient space $\mathbb{R}^{\mathcal{A}_i}$ of each player's strategy space $\mathcal{X}_i = \Delta(\mathcal{A}_i)$, and define the norm $\|\cdot\|$—and the corresponding dual norm $\|\cdot\|_*$—in the joint ambient space $\prod_{j\in\mathcal{N}}\mathbb{R}^{\mathcal{A}_j}$ by $\|x\|^2 := \sum_{i\in\mathcal{N}}\|x_i\|_i^2$, for all $x \in \prod_{j\in\mathcal{N}}\mathbb{R}^{\mathcal{A}_j}$. By multilinearity, each player's individual payoff field v_i is Lipschitz-continuous on $\mathcal{X} = \prod_{j\in\mathcal{N}}\mathcal{X}_j$; we will write L_i for its Lipschitz modulus, that is

$$\|v_i(x) - v_i(x')\|_{i,*} \le L_i\|x - x'\| \quad \text{for all } x, x' \in \mathcal{X}. \tag{D.9}$$

To simplify notation, we will from now on omit the index i from the norm $\|\cdot\|_i$, relying on the context to resolve any ambiguities. For all $x, x' \in \mathcal{X}$, Eq. (D.9) implies that $\sum_{i\in\mathcal{N}}\|x - x'\|^2 = N\|x - x'\|^2 \ge \sum_{i\in\mathcal{N}} 1/L_i^2\|v_i(x) - v_i(x')\|_*^2$, hence

$$\sum_{i\in\mathcal{N}}\|x_i - x_i'\|^2 = \|x - x'\|^2 \ge \sum_{i\in\mathcal{N}}\frac{1}{NL_i^2}\|v_i(x) - v_i(x')\|_*^2 \quad \text{for all } x, x' \in \mathcal{X}. \tag{D.10}$$

Suppose now that the game is harmonic relative to the measure $\mu = (\mu_1, \dots, \mu_N)$, and let $m_i = \sum_{\alpha_i \in \mathcal{A}_i} \mu_{i\alpha_i}$ denote the mass of μ_i. Multiplying both sides of Eq. (D.7) by m_i, summing over every $i \in \mathcal{N}$, and combining Eqs. (D.8) and (D.10), yields

$$\begin{aligned}
&\sum_{t=1}^{T} \sum_{i \in \mathcal{N}} m_i \langle v_{i,t+1/2}, p_i - x_{i,t+1/2} \rangle && \text{(D.11)} \\
&\leq \sum_{i \in \mathcal{N}} m_i \left[\frac{H_i}{\gamma_i \eta_{i,T+1}} + \frac{\gamma_i \|g_{i,1}\|_*^2}{\kappa_i} \right] \\
&+ \sum_{i \in \mathcal{N}} \sum_{t=2}^{T} \left[\frac{m_i \gamma_i \eta_{i,t}}{\kappa_i} \|g_{i,t}\|_*^2 - \frac{\underline{\kappa m}}{4NL_i^2} \left(\|v_{i,t+1/2} - v_{i,t}\|_*^2 + \|v_{i,t} - v_{i,t-1/2}\|_*^2 \right) \right] && \text{\% Eq. (D.10)} \\
&\leq \sum_{i \in \mathcal{N}} \left[\frac{m_i \gamma_i \|g_{i,1}\|_*^2}{\kappa_i} + \frac{m_i H_i}{\gamma_i \eta_{i,T+1}} + \sum_{t=2}^{T} \frac{\gamma_i m_i \eta_{i,t}}{\kappa_i} \|g_{i,t}\|_*^2 - \sum_{t=2}^{T} \frac{\underline{\kappa m}}{8NL_i^2} \|g_{i,t}\|_*^2 \right], && \text{\% Eq. (D.8)}
\end{aligned}$$

with $\underline{m} := \min_{i \in \mathcal{N}} \{m_i\}$ and $\underline{\kappa} := \min_{i \in \mathcal{N}} \{\kappa_i\}$. For the following analysis, it is convenient to define $\sigma_i := \underline{\kappa m} / (16 N L_i^2)$, and rewrite the summand between square brackets in the last line of Eq. (D.11) as

$$[\dots]_i \equiv \frac{m_i \gamma_i \|g_{i,1}\|_*^2}{\kappa_i} \tag{D.12a}$$

$$+ \frac{m_i H_i}{\gamma_i \eta_{i,T+1}} - \sigma_i \sum_{t=2}^{T} \|g_{i,t}\|_*^2 \tag{D.12b}$$

$$+ \sum_{t=2}^{T} \left(\frac{\gamma_i m_i}{\kappa_i} \eta_{i,t} - \sigma_i \right) \|g_{i,t}\|_*^2 . \tag{D.12c}$$

Denote now by $q \in \operatorname{ri} \mathcal{X}$ the strategic center of the harmonic game at hand: if we instantiate Eq. (D.11) at $p \leftarrow q$, its left-hand-side will vanish by Eq. (A.1), leaving us with

$$0 \leq \sum_{i \in \mathcal{N}} \text{(D.12a)}_i + \sum_{i \in \mathcal{N}} \text{(D.12b)}_i(T) + \sum_{i \in \mathcal{N}} \text{(D.12c)}_i(T) \quad \text{for all } T \in \mathbb{N}, \tag{D.13}$$

where we added for clarity the subscript i and the argument (T) to the equation labels.

Note that, up to this point, we have only used the fact that $\eta_{i,t}$ is bounded above—specifically, that $\eta_{i,t} \leq 1$. Beyond this, we have not invoked any particular policy linking the learning-rate sequence to the structure of the difference of gradients $g_{i,t}$.

We now distinguish between the cases *(a)* and *(b)* in the statement of the theorem.

By assumption *(a)*, $\eta_{i,t}$ tautologically converges to a positive constant—its infimum;[7] what remains to be shown is that the sequence of difference of gradients squared $\|g_{i,t}\|_*^2$ is summable.

If assumption *(b)* holds true, by virtue of the adaptive policy (12), $\eta_{i,t}$ converges to a positive limit if and only if $\sum_{t=1}^{\infty} \|g_{i,t}\|_*^2 < \infty$; we have to show that this is indeed the case.

Case *(a)*: meta-policy We claim that, under the assumption *(a)*, the series $\sum_{t=1}^{\infty} \|g_{i,t}\|_*^2$ generated by (FTRL+) in any harmonic game is summable. Indeed, (D.12a)_i does not depend on T, and by the assumption on γ_i it follows that $\text{(D.12c)}_i(T)$ is bounded from above for all $i \in \mathcal{N}$. Thus, invoking Eq. (D.13), we obtain that $\sum_{i \in \mathcal{N}} \text{(D.12b)}_i(T) \geq \mathcal{O}(1)$ for all $T \in \mathbb{N}$. Since $\eta_{i,t}$ is monotonically decreasing with positive infimum, this in turn implies that

$$\sum_{i \in \mathcal{N}} \sigma_i \sum_{t=2}^{T} \|g_{i,t}\|_*^2 \leq \sum_{i \in \mathcal{N}} \frac{m_i H_i}{\gamma_i \eta_{i,\infty}} + \mathcal{O}(1) = \mathcal{O}(1) \quad \text{for all } T \in \mathbb{N}, \tag{D.14}$$

which implies that $\sum_{t=1}^{\infty} \|g_{i,t}\|_*^2 < \infty$ for all $i \in \mathcal{N}$, and we are done. ∎

Remark 6. As discussed in the main text, the meta-policy of case *(a)* mimics the behavior of a constant learning rate, which is positive and "small enough"—where the relevant threshold, as in

[7] The sequence $(\eta_{i,t})_{t \in \mathbb{N}}$ is bounded (by 1 from above and by 0 from below) and non-increasing, hence it converges to its infimum $\eta_{i,\infty}$, which is positive by assumption *(a)*.

assumption Eq. (11), depends on game-specific global quantities that are not available to the players (e.g., harmonic masses, Lipschitz constants of the payoff fields, and the number of players). As we show next, the policy (12) adapts automatically to the observed signals sequence and converges to a positive value *without* requiring access to global information.

Case *(b)*: adaptive policy For each $i \in \mathcal{N}$, assume that $(\eta_{i,t})_{t\in\mathbb{N}}$ follows Eq. (12), and set without loss of generality $\gamma_i = 1$. We then claim that both $(\text{D.12b})_i$ and $(\text{D.12c})_i$ are bounded from above for all $i \in \mathcal{N}$.

Bounding (D.12b). By Eq. (12),

$$\begin{aligned}(\text{D.12b})_i(T) &= m_i H_i \sqrt{1 + \|g_{i,1}\|_*^2 + \sum_{t=2}^{T} \|g_{i,t}\|_*^2} - \sigma_i \sum_{t=2}^{T} \|g_{i,t}\|_*^2 \\ &\leq m_i H_i \sqrt{1 + \|g_{i,1}\|_*^2} + m_i H_i \sqrt{\sum_{t=2}^{T} \|g_{i,t}\|_*^2} - \sigma_i \sum_{t=2}^{T} \|g_{i,t}\|_*^2 \qquad \% \ \sqrt{a+b} \leq \sqrt{a} + \sqrt{b} \\ &\equiv f_i\left(\sqrt{\sum_{t=2}^{T} \|g_{i,t}\|_*^2}\right) \leq \max_{a\in\mathbb{R}} f_i(a) < \infty \quad \text{for all } i \in \mathcal{N} \text{ and } T \in \mathbb{N}\,, \end{aligned} \tag{D.15}$$

where the function $f_i : \mathbb{R} \to \mathbb{R}$, defined as $f_i(a) \coloneqq m_i H_i \sqrt{1 + \|g_{i,1}\|_*^2} + m_i H_i\, a - \sigma_i\, a^2$, is a quadratic function with negative leading coefficient, and as such bounded from above. ◇

Bounding (D.12c). The sequence $(\eta_{i,t})_{t\in\mathbb{N}}$ is bounded (by 1 from above and by 0 from below) and non-increasing, hence it converges to its infimum $\eta_{i,\infty}$, which can be either positive or zero.

- If $\eta_{i,\infty} > 0$, then by Eq. (12) we have $\sum_{t=1}^{\infty} \|g_{i,t}\|_*^2 = 1/\eta_{i,\infty}^2 - 1 < \infty$. Since $\eta_{i,t} \leq 1$,

$$(\text{D.12c})_i(T) \leq \left(\frac{m_i}{\kappa_i} - \sigma_i\right) \sum_{t=2}^{\infty} \|g_{i,t}\|_*^2 < \infty, \quad \text{for all } i \in \mathcal{N} \text{ and } T \in \mathbb{N}\,. \tag{D.16}$$

- If $\eta_{i,\infty} = 0$, then the coefficient multiplying $\|g_{i,t}\|_*^2$ in Eq. (D.12c) becomes negative after finitely many steps. Let $t_i' \in \mathbb{N}$ be the crossing time, namely $t_i' \coloneqq \min_t \{t \in \mathbb{N} : \eta_{i,t}\, m_i/\kappa_i - \sigma_i \leq 0\}$. Then

$$(\text{D.12c}) \leq \sum_{t=2}^{t_i'-1} \left(\frac{m_i}{\kappa_i}\eta_{i,t} - \sigma_i\right) \|g_{i,t}\|_*^2 < \infty, \quad \text{for all } i \in \mathcal{N}\,. \tag{D.17}$$

In either case, Eq. (D.12c) is bounded from above. ◇

Using the facts that $(\text{D.12a})_i$ does not depend on T and that $(\text{D.12c})_i(T)$ is bounded from above, and invoking Eq. (D.13), we obtain that $\sum_{i\in\mathcal{N}} (\text{D.12b})_i$ is bounded from below. Furthermore, since $(\text{D.12b})_i$ is bounded from above for all $i \in \mathcal{N}$, we conclude by Lemma D.3 that $(\text{D.12b})_i$ is also bounded from *below*, for all $i \in \mathcal{N}$. By Eq. (D.15), this implies that $f_i\left(\sqrt{\sum_{t=2}^{T} \|g_{i,t}\|_*^2}\right)$ is bounded from below. Since $\lim_{a\to\infty} f_i(a) = -\infty$, this implies that $\sum_{t=1}^{\infty} \|g_{i,t}\|_*^2 < \infty$; by Eq. (12), this is equivalent to $\eta_{i,\infty} = \lim_{t\to\infty} \eta_{i,t} > 0$. ■

With this result at hand, we are in a position to prove the second part of Theorem 4, which we restate below for convenience:

Theorem 4 (Individual regret guarantees)**.** *Suppose that player $i \in \mathcal{N}$ is running* (FTRL+) *with payoff signals as per* (6)*, parameters as per* (10) *and* (12)*, and $\gamma_i = 1$. Then, letting $H_i \coloneqq \max h_i - \min h_i$ and $G_i = \sup_{x\in\mathcal{X}} \|v_i(x)\|_*$, we have:*

(a) The regret of player $i \in \mathcal{N}$ against any sequence $X_{-i,t} \in \mathcal{X}_{-i}$, $t = 1, 2, \ldots$, is bounded as

$$\mathrm{Reg}_i(T) \leq (H_i + 1/\kappa_i)(1 + 4G_i^2 T)^{1/2} + 2\,G_i^2/\kappa_i = \mathcal{O}\big(\sqrt{T}\big)\,. \tag{14a}$$

(b) If, in addition, all players follow (FTRL+) *as above and the game is harmonic, then*

$$\mathrm{Reg}_i(T) = \mathcal{O}(H_i + 1/\kappa_i)\,. \tag{14b}$$

Proof of Theorem 4–(b). For each $i \in \mathcal{N}$ and any $p_i \in \mathcal{X}_i$, consider the template inequality of Proposition 1 under the choice of parameters of Eq. (10). Sum over $t = 1, 2, \dots, T$, telescope, and invoke Eqs. (B.17) and (B.16c) (the Fenchel-Young inequality and strong convexity) to obtain the bound (D.3) for the individual regret of each player using (FTRL+):

$$\mathrm{Reg}_i(T) = \sum_{t=1}^{T} \langle v_{i,t+1/2}, p_i - x_{i,t+1/2} \rangle \leq \frac{H_i}{\eta_{i,T+1}} + \frac{1}{2\kappa_i} \sum_{t=1}^{T} \eta_{i,t} \|g_{i,t}\|_*^2\,. \tag{D.18}$$

The conclusion follows immediately by Proposition D.1-(b): for all $i \in \mathcal{N}$ and all $T \in \mathbb{N}$, we have that $\eta_{i,T+1} \geq \lim_{t\to\infty} \eta_{i,t} = \eta_{i,\infty} > 0$, and $\sum_{t=1}^{T} \|g_{i,t}\|_*^2 \leq \sum_{t=1}^{\infty} \|g_{i,t}\|_*^2 < \infty$. ∎

We conclude this appendix with two standard technical lemmas used in the proofs of Theorem 4 and Proposition D.1.

Lemma D.2. *Given reals $a_1, \dots, a_T$ such that $\sum_{s=1}^{t} a_s > 0$ for all $t = 1 \dots T$, we have*

$$\sum_{t=1}^{T} a_t \frac{1}{\sqrt{\sum_{s=1}^{t} a_s}} \leq 2\sqrt{\sum_{t=1}^{T} a_t}\,. \tag{D.19}$$

Proof. See Auer et al. [5, Lemma 3.5]. ∎

Lemma D.3. *For any $N \in \mathbb{N}$, let $a\colon (i,t) \in \mathcal{N} \times \mathbb{N} \mapsto a_{i,t} \in \mathbb{R}$ be a finite family of real sequences, with $\mathcal{N} \coloneqq \{1, \dots, N\}$. If the sequence $\left(a_{i,t}\right)_{t\in\mathbb{N}}$ is bounded from above for each $i \in \mathcal{N}$, and the sequence $\left(\sum_{i=1}^{N} a_{i,t}\right)_{t\in\mathbb{N}}$ is bounded from below, then the sequence $\left(a_{i,t}\right)_{t\in\mathbb{N}}$ is bounded from below for each $i \in \mathcal{N}$.*

Proof. By assumption there exist $\underline{C} \in \mathbb{R}$ and $\overline{C}_i \in \mathbb{R}$ such that $\sum_{i\in\mathcal{N}} a_{i,t} \geq \underline{C}$ for all $t \in \mathbb{N}$ and $a_{i,t} \leq \overline{C}_i$ for all $i \in \mathcal{N}$ and $t \in \mathbb{N}$. Then, for any $j \in \mathcal{N}$ and all $t \in \mathbb{N}$, we have

$$a_{j,t} \geq \underline{C} - \sum_{i\neq j} a_{i,t} \geq \underline{C} - \sum_{i\neq j} \overline{C}_i > -\infty\,, \tag{D.20}$$

and we are done. ∎

E Convergence analysis

E.1. Preliminary results.

Proposition E.1 (Convergence of energy levels)**.** *Consider a harmonic game with measure $\mu = (\mu_1, \dots, \mu_N)$ and strategic center $q \in \mathrm{ri}\,\mathcal{X}$, and let $m_i = \sum_{\alpha_i \in \mathcal{A}_i} \mu_{i\alpha_i}$ denote the mass of μ_i. Suppose that each player is running* (FTRL+) *with payoff signals as per Eq. (6) and parameters as per Eq. (10). Assume that either condition 1 or 2 of Proposition D.1 holds. Then, the* energy sequence

$$E_t \coloneqq \sum_{i\in\mathcal{N}} \frac{m_i}{\gamma_i} E_{i,t} \equiv \sum_{i\in\mathcal{N}} \frac{m_i}{\gamma_i \eta_{i,t}} F_i(q_i, \eta_{i,t} y_{i,t}) \tag{E.1}$$

relative to the game's harmonic center converges to some finite value E_∞.

Proof. For each $i \in \mathcal{N}$ and any $p_i \in \mathcal{X}_i$, consider the template inequality of Proposition 1 under the choice of parameters of Eq. (10). By the Fenchel-Young inequality (B.17) and strong convexity (Eq. (B.16c)) we readily obtain

$$\langle v_{i,t+1/2}, p_i - x_{i,t+1/2} \rangle \leq \frac{E_{i,t} - E_{i,t+1}}{\gamma_i} + \left(\frac{1}{\eta_{i,t+1}} - \frac{1}{\eta_{i,t}}\right) \frac{h_i(p_i)}{\gamma_i} + \frac{\gamma_i \eta_{i,t}}{2\kappa_i} \|g_{i,t}\|_*^2\,. \tag{E.2}$$

Instantiating Eq. (E.2) at $p \leftarrow q$, multiplying both sides by m_i, summing over every $i \in \mathcal{N}$, and invoking Eq. (A.1) makes the left-hand-side of Eq. (E.2) vanish. Recalling that $\eta_{i,t} \leq 1$ for all $i \in \mathcal{N}$ and $t \in \mathbb{N}$, we then get

$$E_{t+1} \leq E_t + \epsilon_t \,, \tag{E.3}$$

where

$$\epsilon_t \coloneqq \sum_{i \in \mathcal{N}} m_i \left[\left(\frac{1}{\eta_{i,t+1}} - \frac{1}{\eta_{i,t}} \right) \frac{h_i(q_i)}{\gamma_i} + \frac{\gamma_i \eta_{i,t}}{2\kappa_i} \|g_{i,t}\|_*^2 \right] \tag{E.4}$$

is summable by Proposition D.1. The result follows by a standard property of quasi-Fejér sequences [15, Lemma 3.1]: letting $E'_t = E_t + \sum_{s=t}^{\infty} \epsilon_s$, we have

$$E'_{t+1} = E_{t+1} + \sum_{s=t+1}^{\infty} \epsilon_s \leq E_t + \sum_{s=t}^{\infty} \epsilon_s = E'_t \tag{E.5}$$

by Eq. (E.3). Since E'_t is positive, it converges; and since ϵ_t is summable, so does E_t. ∎

Proposition E.2 (Summability of primal sequences)**.** *Suppose that each player is running* (FTRL+) *in a harmonic game with payoff signals as per Eq. (6) and parameters as per Eq. (10). Assume that either condition 1 or 2 of Proposition D.1 holds. Then, the sequences* $A_t \coloneqq \|x_{t+1/2} - x_t\|^2$ *and* $B_t \coloneqq \|x_{t+1} - x_{t+1/2}\|^2$ *are both summable.*

Proof. For each $i \in \mathcal{N}$ and any $p_i \in \mathcal{X}_i$, consider the template inequality of Proposition 1 under the choice of parameters of Eq. (10). Sum over $t = 1, 2, \dots, T$ and follow the same steps leading to Eq. (D.7)—without extracting terms from the series, nor relabeling indices—to obtain the following bound for sequences generated by (FTRL+):

$$\begin{aligned} \sum_{t=1}^{T} \langle v_{i,t+1/2}, p_i - x_{i,t+1/2} \rangle \leq \frac{H_i}{\gamma_i \eta_{i,T+1}} + \sum_{t=1}^{T} \frac{\gamma_i \eta_{i,t}}{\kappa_i} \|g_{i,t}\|_*^2 \\ - \frac{\kappa_i}{4\gamma_i} \sum_{t=1}^{T} \left[\|x_{i,t+1} - x_{i,t+1/2}\|^2 + \|x_{i,t+1/2} - x_{i,t}\|^2 \right] . \end{aligned} \tag{E.6}$$

Suppose now that the game is harmonic with measure $\mu = (\mu_1, \dots, \mu_N)$ and strategic center $q \in \operatorname{ri} \mathcal{X}$, and let $m_i = \sum_{\alpha_i \in \mathcal{A}_i} \mu_{i\alpha_i}$ denote the mass of μ_i. Similar to Eq. (D.11), instantiating Eq. (E.6) at $p \leftarrow q$, multiplying both sides by m_i, summing over every $i \in \mathcal{N}$, and invoking Eq. (A.1) makes the left-hand-side of Eq. (E.6) vanish. Recalling that $\eta_{i,t} \leq 1$ for all $i \in \mathcal{N}$ and $t \in \mathbb{N}$, and invoking Proposition D.1, one then gets

$$\begin{aligned} \sum_{t=1}^{T} \|x_{t+1} - x_{t+1/2}\|^2 + \sum_{t=1}^{T} \|x_{t+1/2} - x_t\|^2 &\leq \frac{4 \max_{i \in \mathcal{N}} \gamma_i}{\underline{m\kappa}} \sum_{i \in \mathcal{N}} m_i \left[\frac{H_i}{\gamma_i \eta_{i,T+1}} + \sum_{t=1}^{T} \frac{\gamma_i \eta_{i,t}}{\kappa_i} \|g_{i,t}\|_*^2 \right] \\ &\leq \frac{4 \max_{i \in \mathcal{N}} \gamma_i}{\underline{m\kappa}} \sum_{i \in \mathcal{N}} m_i \left[\frac{H_i}{\gamma_i \eta_{i,\infty}} + \frac{\gamma_i}{\kappa_i} \sum_{t=1}^{\infty} \|g_{i,t}\|_*^2 \right] = \mathcal{O}(1) . \end{aligned} \tag{E.7}$$

In particular, the sequences $A_t \coloneqq \|x_{t+1/2} - x_t\|^2$ and $B_t \coloneqq \|x_{t+1} - x_{t+1/2}\|^2$ are both summable. ∎

Corollary 2 (Convergence of offset subsequences)**.** *Under the assumptions of Proposition E.2, if* $x^* \in \mathcal{X}$ *is an accumulation point for the sequence* x_t*, then for any subsequence* x_{t_s} *converging to* x^* *and for any fixed* $r \in \mathbb{N}$*, the* offset subsequences x_{t_s+r} *and* $x_{t_s+r+1/2}$ *converge to* x^**:*

$$\lim_{s \to \infty} x_{t_s+r} = x^* \quad \text{and} \quad \lim_{s \to \infty} x_{t_s+r+1/2} = x^* \,. \tag{E.8}$$

Proof. Since x^* is an accumulation point, x_t admits a subsequence x_{t_s} such that $\lim_{s \to \infty} x_{t_s} = x^*$. Proposition E.2 then implies that also $\lim_{s \to \infty} x_{t_s+1/2} = x^*$. By the triangle inequality, Proposition E.2 implies also that $\lim_{t \to \infty} \|x_{t+1} - x_t\| = 0$ and $\lim_{t \to \infty} \|x_{t+3/2} - x_{t+1/2}\| = 0$, and the conclusion follows by induction. ∎

E.2. Quotient dual space. In this appendix, we present standard techniques employed in Appendix E.3 for quotienting out the degeneracy arising from the fact that the mirror map $Q\colon \mathcal{Y} \to \mathcal{C}$ in Eq. (B.7d) is not injective. We begin by considering a generic convex set $\mathcal{C}$, as in Appendix B, and later specialize to the game-theoretic setting by letting $\mathcal{C} \leftarrow \mathcal{X}$. We refer the reader to [24, App. C.4] for further details.

With the notation of Eq. (B.1), consider the *quotient space* $\mathcal{V}^*/T^\circ$: this is the space of equivalence classes $[y]$ for the equivalence relation on $\mathcal{V}^*$ defined by $y \sim y'$ *if and only if* $y - y' \in T^\circ$, for all $y, y' \in \mathcal{V}^*$. Denote by $\pi\colon \mathcal{V}^* \to \mathcal{V}^*/T^\circ$ the canonical projection onto the quotient space. A key property of the mirror map $Q\colon \mathcal{V}^* \to \mathcal{C}$ is that it *passes to the quotient*, in the sense that it is invariant on the level sets of π. Moreover, by definition, equivalent dual vectors act identically on vectors tangent to $\mathcal{C}$. We collect these two elementary but useful facts in the following

Lemma E.1. *Given the quotient map $\pi\colon \mathcal{V}^* \to \mathcal{V}^*/T^\circ$, the following properties hold:*

$$\text{a. } \textit{If } \pi(y) = \pi(y') \textit{ then } Q(y) = Q(y'), \textit{ for any } y, y' \in \mathcal{V}^*. \tag{E.9a}$$

$$\text{b. } \langle \pi(y), p - x\rangle = \langle y, p - x\rangle \;\; \textit{for all } y \in \mathcal{V}^* \textit{ and all } p, x \in \mathcal{C}. \tag{E.9b}$$

Proof. For the first statement, note that $\pi(y) = \pi(y')$ if and only if $y' = y + n$ for some $n \in T^\circ$, and conclude by Eq. (B.8b). The second statement is immediate by definition and linearity. ∎

Eq. (E.9a) allows to define the *quotient mirror map* $\tilde{Q} \coloneqq Q/_\sim\colon \mathcal{V}^*/T^\circ \to \mathcal{C}$ by $\tilde{Q}([y]) = Q(y)$, for any $y \in \mathcal{V}^*$.

Coordinate system for the quotient space. We now specialize the construction of the previous section to the case where $\mathcal{C} \leftarrow \mathcal{X}$ is the standard simplex in $\mathcal{V} \coloneqq \mathbb{R}^{K+1}$. In this case $\operatorname{aff}\mathcal{X} = \{x \in \mathcal{V} : \sum_{\alpha=0}^K x_\alpha = 1\}$, $T = \{x \in \mathcal{V} : \sum_{\alpha=0}^K x_\alpha = 0\}$, and $T^\circ = \operatorname{span}(1) \subset \mathcal{Y} \coloneqq \mathcal{V}^*$. Introducing a coordinate system on $\mathcal{V}^*/T^\circ$ amounts to selecting, for each equivalence class $[y]$, a representative element. For any $y \in \mathcal{V}^*$, a convenient choice is the unique element $z \in [y]$ that intersects, for instance, the hyperplane $\mathcal{Z} \coloneqq \{y \in \mathcal{Y} : y_0 = 0\}$. This yields an isomorphism $\mathcal{V}^*/T^\circ \cong \mathcal{Z} \subset \mathcal{Y}$ via the mapping $[y] \mapsto z \coloneqq (0, y_1' - y_0', \dots, y_K' - y_0')$, which is well defined for any $y' \in [y]$. Following a standard convention, we henceforth identify $\mathcal{V}^*/T^\circ$ with $\mathcal{Z}$, and write the canonical projection π and the quotient mirror map $\tilde{Q}$ as

$$\pi\colon y \in \mathcal{Y} \mapsto z = (0, y_1 - y_0, \dots, y_K - y_0) \in \mathcal{Z}\,, \tag{E.10a}$$

$$\tilde{Q}\colon z \in \mathcal{Z} \mapsto \tilde{Q}(z) = Q(y) \in \mathcal{X} \quad \text{for any } y \in \pi^{-1}(z)\,. \tag{E.10b}$$

We conclude this section with a technical lemma used in the proof of Theorems 2 and 3.

Lemma E.2. *Let y_t be a sequence in $\mathcal{Y}$, and let p be a point in the relative interior of the simplex $\mathcal{X}$. If $\sup_{t\in\mathbb{N}} |h^*(y_t) - \langle y_t, p\rangle| < \infty$, then the sequence $z_t = \pi(y_t)$ is bounded in $\mathcal{Z}$.*

Proof. See [28, Appendix D]. ∎

E.3. Convergence to the set of Nash equilibria. With all this in hand, we are finally in a position to prove—simultaneously, by virtue of Proposition D.1—our main convergence results:

Theorem 2 (Convergence to equilibrium)**.** *Suppose that each player is running* (FTRL+) *in a harmonic game with payoff signals as per* (6) *and parameters as per* (10)*. If $\liminf_t \eta_{i,t} > 0$ and $\gamma_i > 0$ is sufficiently small, then x_t converges to the set of Nash equilibria of the game.*

Theorem 3 (Adaptive convergence)**.** *Suppose that each player is running* (FTRL+) *in a harmonic game with payoff signals as per* (6)*, parameters as per* (10) *and* (12)*, and $\gamma_i = 1$. Then x_t converges to the set of Nash equilibria of the game.*

Remark 7. In the following we invoke standard techniques to quotient out the degeneracy arising from the fact that the mirror map $Q\colon \mathcal{Y} \to \mathcal{X}$ of Eq. (B.7d) is not injective; the relevant details are provided in Appendix E.2.

Proof of Theorems 2 and 3. Our proof proceeds in two steps:

Step 1: x_t converges to set of accumulation points. Convergence of the energy sequence—Proposition E.1—implies that x_t converges to $\tilde{Q}(\Omega_{z_t}) \subseteq \Omega_{x_t} \cap \mathcal{X}_h$, where Ω_{x_t} is the

set of accumulation points of x_t; $\tilde{Q}(\Omega_{z_t})$ is the image via the quotient mirror map (E.10b) of the set of accumulation points Ω_{z_t} of the sequence z_t, where $z_{i,t} \coloneqq \pi_i(\eta_{i,t} y_{i,t})$; π is the canonical projection of (E.10a); and $\mathcal{X}_h$ is the domain of subdifferentiability (or prox-domain) of h on $\mathcal{X}$.

Step 2: Accumulation points are Nash equilibria. Convergence of the offset subsequences—Corollary 2—implies that $\tilde{Q}(\Omega_{z_t}) \subseteq \mathrm{NE}$, where NE is the set of Nash equilibria of the harmonic game.

We detail these steps below.

Step 1: x_t converges to set of accumulation points. We now proceed to show that x_t converges to $\tilde{Q}(\Omega_{z_t}) \subseteq \Omega_{x_t} \cap \mathcal{X}_h$. By Proposition E.1, $E_t = \sum_{i\in\mathcal{N}}(m_i/\gamma_i)E_{i,t}$ converges; in particular, $F_i(q_i, \eta_{i,t}y_{i,t}) = h_i(q_i) + h_i^*(\eta_{i,t}y_{i,t}) - \eta_{i,t}\langle y_{i,t}, q_i\rangle$ is bounded from above for each $i \in \mathcal{N}$. By Lemma E.2, this implies that each component of $z_{i,t} = \pi_i(\eta_{i,t}y_{i,t})$ is bounded, hence the sequence z_t itself is bounded in $\mathcal{Z} = \prod_{j\in\mathcal{N}} \mathcal{Z}_j$. By Lemma E.3, this implies that z_t converges to Ω_{z_t}, the set of its accumulation points. By continuity of the mirror map—Item 4 in Proposition B.1—the sequence x_t, with $x_{i,t} = \tilde{Q}_i(z_{i,t}) = Q_i(\eta_{i,t}y_{i,t})$, then converges to $\tilde{Q}(\Omega_{z_t})$. Clearly, $\tilde{Q}(\Omega_{z_t}) \subseteq \Omega_{x_t} \cap \mathcal{X}_h$: if $x^* \in \tilde{Q}(\Omega_{z_t})$ then by continuity $x^* = \tilde{Q}(z^*) = \tilde{Q}(\lim_{s\to\infty} z_{t_s}) = \lim_{s\to\infty} \tilde{Q}(z_{t_s}) = \lim_{s\to\infty} x_{t_s}$, for some convergent subsequence, showing that $x^* \in \Omega_{x_t}$; and $\operatorname{im}\tilde{Q} \subseteq \operatorname{im} Q = \mathcal{X}_h$, the prox-domain of h on $\mathcal{X}$.

Step 2: Accumulation points are Nash equilibria. We now proceed to show that every point x^* in $\tilde{Q}(\Omega_{z_t})$—which is non-empty, by the previous step—is a solution to (SVI), and hence a Nash equilibrium. For any $x^* = \tilde{Q}(z^*) \in \tilde{Q}(\Omega_{z_t})$, let z_{t_s} be a subsequence of z_t converging to z^*; then $x_{t_s} \to x^*$ by continuity. For any fixed $r \in \mathbb{N}$, a straightforward induction on the update rule of (FTRL+) shows that

$$y^i_{t_s+r} = y^i_{t_s} + \gamma_i \sum_{j=0}^{r-1} v^i_{t_s+1/2+j}\,, \tag{E.11}$$

where we promoted the player index $i \in \mathcal{N}$ to a superscript to avoid clutter. Invoking Eq. (B.13) for any $p \in \mathcal{X}$ with $y' \leftarrow \eta^i_{t_s+r} y^i_{t_s+r}$ and $y \leftarrow \eta^i_{t_s+r} y^i_{t_s}$ yields

$$\begin{aligned}
\left\langle \eta^i_{t_s+r} y^i_{t_s+r} - \eta^i_{t_s+r} y^i_{t_s}, p_i - Q_i(\eta^i_{t_s+r} y^i_{t_s+r})\right\rangle &= \eta^i_{t_s+r}\gamma_i \left\langle \sum_{j=0}^{r-1} v^i_{t_s+1/2+j}, p_i - x^i_{t_s+r}\right\rangle \\
&\le \left\langle \nabla h_i(x^i_{t_s+r}) - \eta^i_{t_s+r} y^i_{t_s}, p_i - x^i_{t_s+r}\right\rangle \\
&= \left\langle \nabla h_i(x^i_{t_s+r}) - \frac{\eta^i_{t_s+r}}{\eta^i_{t_s}} z^i_{t_s}, p_i - x^i_{t_s+r}\right\rangle. \quad \%\ \text{Eq. (E.9b)}
\end{aligned} \tag{E.12}$$

Recall that $v^i_{t_s+1/2+j} \equiv v_i(x_{t_s+1/2+j})$. Passing to the limit $s \to \infty$, by Corollary 2, Proposition D.1, and the continuity of the subgradient selection,

$$\eta_{i,\infty}\gamma_i \left\langle \sum_{j=0}^{r-1} v_i(x^*), p_i - x^*_i\right\rangle = \eta_{i,\infty}\gamma_i\, r \left\langle v_i(x^*), p_i - x^*_i\right\rangle \le \left\langle \nabla h_i(x^*_i) - z^*_i, p_i - x^*_i\right\rangle = \mathcal{O}(1)\,, \tag{E.13}$$

for all $r \in \mathbb{N}$, all $i \in \mathcal{N}$, and all $p \in \mathcal{X}$. The right-hand side of the inequality does not depend on r, hence if the left-hand side were positive, by choosing $r \in \mathbb{N}$ big enough one would get to a contradiction. Hence, it must hold true that

$$\left\langle v_i(x^*), p_i - x^*_i\right\rangle \le 0 \quad \text{for all } i \in \mathcal{N} \text{ and all } p \in \mathcal{X}\,. \tag{E.14}$$

Hence, by (SVI), x^* is a Nash equilibrium.

In conclusion, we showed that the sequence x_t generated in any harmonic game by FTRL+ under the assumptions of the theorem converges to the set $\tilde{Q}(\Omega_{z_t})$, and that every point in $\tilde{Q}(\Omega_{z_t})$ is a Nash equilibrium for the harmonic game at hand. ∎

We conclude this section with a technical lemma used in the proof of Theorems 2 and 3:

Lemma E.3. *Let z_t be a sequence in a finite-dimensional Euclidean space. If z_t is bounded, then it converges to its set of accumulation points.*

Proof. Every convergent subsequence z_{t_s} of z_t—of which at least one exists by the Bolzano-Weierstrass theorem—converges to Ω, the set of accumulation points of z_t. We want to show that for every open neighborhood $\mathcal{U}$ of Ω, $z_t \in \mathcal{U}$ for all t big enough. Assume by contradiction this is not the case: there exists an open neighborhood $\mathcal{U}$ of Ω and a subsequence z_{t_s} such that $z_{t_s} \notin \mathcal{U}$ for all $s \in \mathbb{N}$. Since z_{t_s} is bounded it admits a convergent subsequence: $\lim_{\tau\to\infty} z_{t_{s_\tau}} = x^* \notin \Omega$, a contradiction. ∎

F Additional numerical details

This appendix gives the implementation details of the experiments in Section 6, reported in Fig. 2.

Games and coordinates. Both games have three players and two actions per player. Their payoff tables are

	Connected Nash set			Disconnected Nash set		
(a_1, a_2, a_3)	u_1	u_2	u_3	u_1	u_2	u_3
$(0,0,0)$	-4	15	-7	-18	3	21
$(0,0,1)$	-4	-2	1	7	16	7
$(0,1,0)$	7	1	-1	1	8	1
$(0,1,1)$	1	-2	-1	3	0	5
$(1,0,0)$	1	1	-1	-9	-3	-5
$(1,0,1)$	0	-2	1	9	-6	7
$(1,1,0)$	0	0	1	2	-6	9
$(1,1,1)$	1	0	-1	-9	8	7

(F.1)

These are the left and middle games in Fig. 1; the regret experiment uses the latter. We write $x_i \in [0,1]$ for the probability of action 1, and $v_i(x) = u_i(1; x_{-i}) - u_i(0; x_{-i})$ for the effective payoff vector. Binary entropy gives the choice map $Q_i(y_i) = (1 + e^{-y_i})^{-1}$. All experiments use the update in (FTRL+) with the signals as in Eq. (6) and $\lambda_i = 1/2$.

Convergence panels. Panels (a,b) of Fig. 2 illustrate Theorem 3. Each game is run for 5,000 rounds with the adaptive policy in Eq. (12) and score step $\gamma_i = 1$. Both games use the same ten initial strategy profiles, with each player's probability of action 1 between 0.05 and 0.95. The starting probabilities are spread across this interval using Latin-hypercube sampling with fixed random seed 20260801. For each initial profile x_1, the scores are $y_{i,1} = \log\big(x_{i,1}/(1 - x_{i,1})\big)$. The plotted quantity is the maximum unilateral Nash gap at the played leading state:

$$\mathrm{Gap}(x) = \max_i \max_{a_i \in \{0,1\}} [u_i(a_i; x_{-i}) - u_i(x)]. \tag{F.2}$$

At each recorded round, we plot its median and 25th–75th percentiles across the ten runs. Rounds are linear and gaps logarithmic.

Regret configurations and initialization. The robustness comparison in panel (c) of Fig. 2 illustrates the individual and self-play guarantees of Theorem 4. We run four configurations for 8,000 rounds: all players use constant score steps $\gamma_i = 0.05$ and learning rates $\eta_{i,t} = 1$; all players use $\gamma_i = 1$ and the adaptive learning rates in Eq. (12); or either policy is used by players $1, 2$ while player 3 instead uses constant $\gamma_3 = 5$ and $\eta_{3,t} = 1$. The adaptive learning rates start at 1, and the initial preceding payoff signal is zero. We use reference scores $\bar{y} = (1, -2, 0.5)$: a constant-policy player starts at $y_{i,1} = \gamma_i \bar{y}_i$, and an adaptive player at $y_{i,1} = \bar{y}_i$. This nonzero initialization differs from the simplifying convention (Init).

Regret measurement and checks. Regret is evaluated on the leading states actually played. In binary effective coordinates it is exactly

$$\mathrm{Reg}_i(T) = \max\left\{0, \sum_{t=1}^{T} v_i(x_{t+1/2})\right\} - \sum_{t=1}^{T} x_{i,t+1/2} v_i(x_{t+1/2}). \tag{F.3}$$

The maximum is taken over all three players in self-play and over players $1, 2$ when player 3 deviates. At the final round, the constant and adaptive deviation runs give approximately 118.55 and 5.46, respectively; their least-squares log–log slopes over rounds 4,001–8,000 are approximately 1.01 and 0.

G Further related work

G.1. Extrapolation in games. Classical extra-gradient and optimistic methods have their roots in saddle-point optimization and variational-inequality theory [23, 31, 36]. For *bilinear two-player zero-sum games*, minimax arguments establish last-iterate convergence of optimistic gradient and multiplicative-weights methods [16, 46]. Beyond the bilinear setting, convergence analyses of Mirror-Prox and optimistic mirror methods typically rely on *monotonicity* of the payoff operator [31]: under our payoff-maximization convention, this means that $\langle v(x) - v(x'), x - x' \rangle \leq 0$ for all $x, x' \in \mathcal{X}$. A related condition is *variational stability*: a nonempty Nash set $\mathcal{X}^\star$ of a continuous game is variationally stable if $\langle v(x), x - x^\star \rangle \leq 0$ for all $x \in \mathcal{X}$ and $x^\star \in \mathcal{X}^\star$. This condition provides a signed drift toward equilibrium, which, under *strict* variational stability or suitable global regularity of the regularizer, can be leveraged to establish convergence of regularized and optimistic learning methods [22, 27]; see Table 1 for a schematic summary.

Crucially, none of these arguments applies to harmonic games: this class does contain bilinear zero-sum games with a fully mixed equilibrium (which are monotone), but it also contains highly non-monotone instances: Theorem 1 shows that every odd-dimensional harmonic game possess an algebraic variety of Nash equilibria of dimension at least one, which is non-convex in general.

Setting	**Structure used**	**Guaranteed by harmonicity?**
Convex–concave saddle problems [31]	Monotonicity	✗
Bilinear two-player zero-sum games [16, 46]	Minimax structure	✗
Variationally stable games [22, 27]	Variational stability	✗

Table 1: Convergence structures used in representative extrapolation results.

G.2. Adaptive methods in games. Adaptive extra-gradient methods have previously been studied in zero-sum and variationally stable games. The seminal work of Rakhlin & Sridharan [38] introduced an adaptive policy guaranteeing *average* convergence to minimax equilibria in zero-sum games, a weaker guarantee than the last-iterate convergence of Theorem 3. More recently, Hsieh et al. [22] developed adaptive extra-gradient schemes achieving last-iterate convergence and constant regret in variationally stable games. Harmonic games, however, do not satisfy variational stability, so their convergence analysis does not carry over to our setting. This limitation is particularly significant for entropy-like regularizers: Theorem 7 of Hsieh et al. [22] requires *strict* variational stability, which entails a consistent drift toward equilibrium and is explicitly violated in harmonic games. Our convergence proof in Appendix E bypasses this requirement through a fundamentally different convergence technique.

G.3. Convergence in harmonic games. As discussed in Section 5, Legacci et al. [24] proved convergence to the Nash set and bounded self-play regret for FTRL+ with constant, fine-tuned parameters. This comes with several limitations:

1. *Game-dependent tuning.* The admissible step size depends on global quantities, including the harmonic weights, which in turn determine an equilibrium of the underlying game (Proposition A.2).
2. *Adversarial fragility.* Bounded regret in self-play does not give a worst-case regret bound when an opponent deviates; linear regret can occur, as illustrated in Section 6.
3. If the players attempt to mitigate this by using a vanishing learning-rate policy, e.g., proportional to $1/\sqrt{t}$, both the $\mathcal{O}(1)$ self-play regret guarantee and the convergence argument would break down, since these require the learning rate to stay bounded away from zero.

Theorems 3 and 4 resolve the above issues by proposing a method which, simultaneously, (*a*) is run with only locally-available information; (*b*) converges to the Nash set; and (*c*) remains robust against adversarial play while retaining a much faster $\mathcal{O}(1)$ regret minimization in self-play.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: our main contributions, Theorems 1–4, are presented the abstract and introduction.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The closing paragraph of Section 7 discusses limitations.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

Justification: The proofs of Theorems 1–4 are detailed in the appendix.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Appendix F details the experimental setup.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [N/A] Justification: The paper contains no data or code.

   Guidelines:

- The answer [N/A] means that the paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://neurips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Appendix F details the experimental setting.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: The convergence plots in Fig. 2 report medians and interquartile bands (25th–75th percentiles) across ten initializations per game; see Appendix F.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).

- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Appendix F details the experimental setting.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: The authors reviewed the NeurIPS Code of Ethics and their research conforms with it.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

    Justification: The results are theoretical in nature and carry little societal impact.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The paper poses no such risks.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [N/A]

    Justification: The paper does not use existing assets.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: the paper does not release new assets.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.

- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: the paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: the paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: LLMs were used for writing, editing, or formatting purposes; for running and plotting numerical experiments; and for literature review. The core results and proofs in this research do not involve LLMs.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.

- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.